A scalable consistent second-order SPH solver for unsteady low Reynolds number flows

Nathaniel Trask, Martin Maxey

Brown University: Division of Applied Mathematics, 182 George St. Providence, RI 02906

Kyungjoo Kim, Mauro Perego, Michael L. Parks

Computational Mathematics, Sandia National Laboratories, P.O. Box 5800, MS 1320, Albuquerque, NM 87185.

Kai Yang, Jinchao Xu

Department of Mathematics, Pennsylvania State University, University Park, PA 16802

Abstract

Smoothed Particle Hydrodynamics (SPH) has successfully been used to study a variety of cases involving nearly inviscid flows where conservation properties allow for good physical approximation despite poor theoretical approximation properties of differential operators. When used to study unsteady low Reynolds number flow with large dissipation, conservation alone cannot ensure quality of approximation and the traditional approach is inconsistent. An alternative formulation has recently become popular making use of an approximate splitting scheme to ensure a divergence-free velocity field. However, this scheme relies on an inconsistent discretization of the Laplacian that diverges as particles become disordered under flow. We present an incremental pressure correction scheme and combination of existing differential operator renormalizations that are able to achieve second order accuracy in time and space. A brief review of SPH approximation theory is provided to highlight the necessity of these renormalizations in implementing an approximate factorization scheme. We demonstrate that when fast algebraic multigrid preconditioners are used to solve the resulting linear systems, the scheme results in a consistent approximation that is scalable and amenable to parallelization. Several validation cases are presented for which a speedup of several orders of magnitude is achieved over traditional SPH approaches. Finally, an interface has been developed between the particle library LAMMPS and the sparse linear algebra libraries in Trilinos providing a massively parallel 3D SPH capability. Scaling results for up to 134 million particles on 32,768 cores are presented along with a demonstration of the capability to simulate complex 3D geometries. These results show that the added complexity of applying the necessary consistency corrections actually provides a factor of four speed-up per linear solver iteration versus the uncorrected case, despite the additional cost of constructing the corrections. The resulting library provides a method that is consistent, efficient, and second order in both space and time while maintaining the flexibility of classical SPH for single phase flows.

Keywords: Smooth particle hydrodynamics, ISPH, splitting schemes, meshless method

1. Introduction

Smoothed Particle Hydrodynamics (SPH) is a meshless particle-based method commonly applied to solve partial differential equations in problems arising in fluid mechanics. A set of moving collocation points (or particles) are distributed throughout a domain, and the Lagrangian form of the Navier-Stokes equations is solved at each collocation point. By applying a kernel approximation to evaluate the weak form of the differential operators in the equations, the process gives a means for developing a set of ordinary differential equations describing the trajectories of the fluid particles as the flow evolves in time. By using a compact kernel, the interaction between particles can be limited to nearest neighbors resulting in sparse differential operators and an efficient method. Oftentimes the collocation points are modeled as small Lagrangian fluid elements of fixed finite mass, providing a physical interpretation of the

approximation process [1]. This interpretation turns the method into a useful modeling platform for a wide variety of problems [2].

The method was originally developed in an astrophysical context for solving the Euler equations [5, 6]. In these problems it is only necessary to accurately approximate first order derivatives arising from the pressure gradient and velocity divergence terms, while the second order viscous operator present in the Navier-Stokes equations is added only artificially for stability reasons to control particle distribution and spurious numerical fluctuations near discontinuities [3]. Further, for these problems it is possible to derive the discrete system of equations from a Lagrangian variational formulation, ensuring that the resulting ODE has a Hamiltonian structure [4, 7]. Theoretical work has revealed a sensitivity of the derivative operator in SPH to the underlying particle configuration [8]. As the particle distribution becomes more anisotropic and departs from a regular configuration, the approximation is only convergent in the limit as both the number of particles in the domain and the number of neighbors in the support of the kernel are simultaneously increased. Because in practice this decreases the efficiency of the method, the quality of the SPH approximation often relies on the conservation properties arising from this Hamiltonian derivation to achieve useful results rather than attempting to achieve an answer that converges to the exact solution [9]. Although inconsistent, for these problems the pressure gradient operator actually indirectly maintains a quasi-uniform particle distribution through the implicit imposition of a background pressure. As a result, the standard formulation of SPH is only quasi-Lagrangian. Subsequent work has shown that by correcting for this background pressure good results can be achieved for a variety of flows [10, 66, 67].

Although the classical SPH operators are inconsistent, they have historically been useful for a class of problems where formal accuracy is of secondary importance. The other area where SPH has shown significant success is in simulating interfacial multiphase flows with large density ratios. While traditional mesh-based methods need to explicitly track the location of a topologically complex interface, the SPH formulation handles these flows naturally and only needs to apply appropriate surface tension forces [11, 12]. The problems where this approach is applied are again largely inviscid and the physics are governed primarily by the accuracy of the first derivative operator and maintaining the conservation properties of the Hamiltonian formulation [13]. In recent years SPH has received attention as a useful platform for studying low Reynolds number flows in microfluidic applications [14, 15, 16]. The meshless nature of SPH allows for the study of fluid structure interaction, such as those found in colloidal suspensions, without the need either to adaptively fit a mesh to the moving boundary or to couple the solid boundary to a fixed Eulerian mesh through an artificial diffuse boundary condition with the immersed boundary methods [17]).

For these low Reynolds number problems the physics are dominated by the viscous operator in the Navier-Stokes equation, and due to this dissipative property the Hamiltonian arguments that have preserved the quality of under-resolved SPH in other applications are no longer applicable. Further, numerical experiments [18, 19] and some analytic estimates [20] have demonstrated that the second order operator is more sensitive to particle anisotropy, and actually diverges as the number of particles used in the domain increases. In the context of other mesh-free methods, it has been known for some time that an approximation satisfying the local polynomial reproduction property (e.g. the highest degree polynomial for which the approximation is exact) can be used to guarantee a consistent approximation [21, 22]. Randles and Libersky [61], Oger et al. [23] and, more recently, Fatehi and Mansari [20] have developed such corrections for the gradient and Laplacian operators, respectively. These methods function by inverting a pair of small symmetric matrices for each particle in the domain constructed in such a way as to guarantee the exact evaluation of each operator when applied to a linear function. Provided that the particles are arranged “well enough” so that such an approximation exists, this gives consistent approximation and the error no longer needs to be controlled by increasing the diameter of the kernel support. Although this introduces additional work to construct and invert these correction tensors, the resulting linear operators require fewer neighbors and therefore are more sparse, requiring for example only 20 neighbors per particle to compute a Laplacian in 2D compared to the 80 particles that would commonly be used in classical SPH [14]. In 3D this effect is more dramatic, requiring for example 90 neighbors to compute a Laplacian compared to 730 using the classical approach.

An additional recent development in SPH is the use projection methods to satisfy the solenoidal constraint on the velocity field in the incompressible Navier-Stokes equations. The classical formulation of SPH applies an artificial compressibility assumption to control the divergence error in the so-called weakly compressible formulations (WC-SPH). By enforcing a linear equation of state relating the density error to the pressure, the sound speed assumes the role of a penalty parameter for which the divergence error can be controlled at the expense of a stiffer CFL condition resulting from corresponding artificial acoustic modes [24]. This approach is consistent only for steady-state flows or

in the limit as the sound speed is increased [67], and at low Reynolds number exchanges the coupled elliptic system for a hyperbolic one requiring appropriate non-reflecting pressure boundary conditions [25]. In practice these limitations are often overlooked and the artificial compressibility assumption yields an explicit method that is simple to code and amenable to exploit fine-grained parallelism on modern computing architectures e.g. GPUs, allowing a fast simulation of large numbers of particles. Recent implementations on GPUs have demonstrated nearly ideal scaling results for up to 10^9 particles in this formulation [26, 27].

Alternatively, Cummins and Rudman applied a projection method (ISPH) where a Poisson problem is solved at every timestep to satisfy the divergence-free constraint to within the accuracy of the Laplacian operator [28]. Due to the fact that SPH effectively uses equal order approximation for the pressure and velocity, Cummins and Rudman noted the presence of a checkerboard instability analogous to that found in cell-centered, mesh based methods or mixed finite element discretizations failing the inf-sup condition [29]. To remove this they employ an approximate factorization [33], using a discrete Laplacian operator in the Poisson solver rather than the composition of the discrete divergence and discrete gradient operators. Because of this, the divergence constraint is only satisfied to within the truncation error of the discrete Laplacian operator, and in light of Fatehi’s analysis [20], the error will diverge when a flow is simulated in high resolution.

While solving the resulting linear systems in ISPH incurs a large increase in computational cost and complexity, it removes the artificial stability constraint imposed in WCSPH and allows for larger timesteps that offset the computational overhead of the method for a large class of problems [30, 31, 32, 34]. Unlike WCSPH, this method is consistent for transient problems, but the large timesteps required to obtain a cost-effective simulation introduce temporal error. Higher order, stiffly stable methods can remedy this and are well-established for projection schemes in grid-based methods [37]. Furthermore, in contrast to the approach used in classical SPH, the viscous term can be treated implicitly to avoid a costly viscous stability restriction at low Reynolds number, and higher order splitting schemes can be used to balance the error introduced by taking large timesteps. When the resulting system is solved using efficient multigrid preconditioners, this technique has proven to be fast, accurate, and scalable for a variety of other discretizations.

In the present work, we present an $O(N)$ algorithm for efficiently computing the correction terms introduced by Fatehi and Mansari. So as to illustrate the necessity of these corrections for solving second order elliptic PDEs, we demonstrate that we are able to maintain second order convergence in space for a set of test problems indicative of typical SPH particle configurations. These corrections produce sparser discrete operators leading ultimately to a less expensive system of linear equations to be solved. We then introduce a second-order, incremental pressure projection scheme modified to be consistent in a Lagrangian formulation. We present near optimal scaling results that we have been able to achieve for the resulting Helmholtz and Poisson problems using a serial algebraic multigrid library.

To obtain a fair assesment of the method’s performance, the standard uncorrected WCSPH and ISPH algorithms were also coded using the same data structures and algorithms and are used to compare accuracy and computational efficiency against the new scheme. The Taylor-Green vortex case is used to show second order convergence of the new method in time. We demonstrate that by combining the consistency corrections, the higher order splitting scheme, and the fast linear algebra solvers we are able to achieve an order of magnitude increase in computational efficiency over classical methods while maintaining scalable results. With the efficiency and stability of the method demonstrated, we then develop a parallel implementation of the algorithm using the massively parallel particle library LAMMPS [58] and the linear algebra library Trilinos [57] to implement the scheme in parallel for 3D problems. We demonstrate that our implementation combining the two libraries is scalable for a series of problems consisting of up to 134 million particles. Flow through a body-centered cubic lattice is used to demonstrate the capability of the method to match the capabilities of standard finite volume results in non-trivial geometries.

2. Modelling approach

2.1. Governing equations

We seek an approximation of the Lagrangian form of the incompressible Navier-Stokes equations:

$$\begin{cases} \frac{d\mathbf{u}}{dt} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{u}, \\ \nabla \cdot \mathbf{u} = 0, \\ \frac{d\mathbf{x}}{dt} = \mathbf{u}, \end{cases} \quad (1)$$

for velocity ($\mathbf{u}$), density (ρ), position ($\mathbf{x}$) and where pressure (p) is a Lagrange multiplier used to satisfy the divergence-free constraint and the density is constant. By solving with the material derivative on the left hand side of the equation we obtain a coupled set of linear PDEs which yield an ODE for the trajectories of each particle. If we naïvely solve for the monolithic coupled system, we obtain a system of equations with $4N$ unknowns, where N is the number of SPH particles. In analogy to the system resulting from simulating compressible gas flows, Chorin [24] proposed replacing the incompressibility constraint with an equation of state relating the pressure to deviations in the density,

$$\begin{cases} \frac{d\mathbf{u}_{ac}}{dt} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{u}_{ac}, \\ \frac{d\rho}{dt} + \nabla \cdot \mathbf{u}_{ac} = 0, \\ p = f(\rho), \end{cases} \quad (2)$$

where we define the velocity resulting from the weakly compressible formulation as $\mathbf{u}_{ac}$. A typical choice for the equation of state is the linear relation $p = c^2(\rho - \rho_0) + p_0$. ρ_0 and p_0 are selected as a suitable reference density and pressure for the flow and c is the artificial sound speed. This speed is typically chosen to be much lower than the physical sound speed of the fluid to avoid an expensive stability restriction but large enough to suppress density fluctuations to a desired level. In SPH applications, this value is frequently set to $c = 10U_{char}$, where U_{char} is a characteristic velocity of the flow. It can be shown that Equation 2 is equivalent to solving the system

$$\begin{cases} \frac{d\mathbf{u}_{ac}}{dt} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{u}_{ac}, \\ \frac{dp}{dt} + c^2\nabla \cdot \mathbf{u}_{ac} = 0. \end{cases}$$

In this form we can see that the sound speed plays the role of a penalty parameter, setting the pressure in such a way as to penalize error in the divergence. This approximation is only consistent with Equation 1 for either steady low Reynolds number flows where $\frac{dp}{dt} = 0$ or for large c . In fact, it is possible to derive an approximation for the modeling error $\mathbf{e} = \mathbf{u} - \mathbf{u}_{ac}$ from the continuity equation

$$\nabla \cdot \mathbf{e} = \nabla \cdot (\mathbf{u} - \mathbf{u}_{ac}) = -\frac{1}{\rho c^2} \frac{dp}{dt} \quad (3)$$

from which it is possible to derive the following lower bound on the error induced by the artificial sound speed.

$$\|\mathbf{e}\|_{H^1} \geq \frac{3}{\rho c^2} \left\| \frac{dp}{dt} \right\|_{L^2} \quad (4)$$

where $\|\cdot\|_{H^1}$ denotes the $W^{1,2}$ Sobolev norm and $\|\cdot\|_{L^2}$ denotes the standard L_2 norm. In the SPH literature, authors frequently study convergence only by increasing the number of particles and fail to investigate whether they are converging to this inconsistent result when studying transient problems.

An alternative approach for decoupling the system of equations also originally proposed by Chorin [38] is to solve Equation 1 without attempting to satisfy the divergence-free constraint. A projection into a divergence-free space is then performed on the initial prediction to obtain a solenoidal velocity at the end of the time step. Applying this correction introduces a splitting error dependent on the size of the timestep. Since its original introduction, the method

has been generalized to achieve arbitrarily high order polynomial reduction of the splitting error and is the standard approach for mesh-based methods for incompressible flows [36].

Cummins and Rudman's adaptation of Chorin's scheme has become the standard approach in SPH under the name incompressible smoothed particle hydrodynamics (ISPH) [28], although more recently an alternative higher order scheme was investigated by Hosseini and Feng [34]. In this approach the semi-discretized Lagrangian Navier-Stokes equations for flow with rigid, fixed boundaries

$$\begin{cases} \frac{\mathbf{u}^{n+1} - \mathbf{u}^n}{\Delta t} = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \mathbf{u}^{n+1} \\ \nabla \cdot \mathbf{u} = 0 \end{cases} \quad (5)$$

are split into the following two predictor/corrector steps:

$$\begin{cases} \frac{\mathbf{u}^* - \mathbf{u}^n}{\Delta t} = \nu \nabla^2 \mathbf{u}^* & x \in \Omega, \\ \mathbf{u}^* = 0 & x \in \partial\Omega, \end{cases} \quad (6)$$

$$\begin{cases} \frac{\mathbf{u}^{n+1} - \mathbf{u}^*}{\Delta t} = -\frac{1}{\rho} \nabla p^{n+1} & x \in \Omega, \\ \nabla \cdot \mathbf{u}^{n+1} = 0 \\ \mathbf{u}^{n+1} \cdot \hat{n} = 0 & x \in \partial\Omega. \end{cases} \quad (7)$$

To determine p^{n+1} with which to correct the initial estimate, the divergence of the second equation gives a Poisson problem for the pressure.

$$\begin{cases} \frac{1}{\rho} \nabla^2 p^{n+1} = \frac{\nabla \cdot \mathbf{u}^*}{\Delta t} & x \in \Omega, \\ \partial_n p^{n+1} = 0 & x \in \partial\Omega. \end{cases} \quad (8)$$

This scheme has been proven to provide $O(\Delta t)$ convergence in velocity and $O(\Delta t^{0.5})$ convergence in pressure for Galerkin finite element schemes. In SPH, the viscous term in the prediction step is often treated explicitly to avoid inverting the resulting Helmholtz matrix. We point out that for most problems, the Poisson matrix is much more poorly conditioned and there is little to be gained by treating the viscous term explicitly; in fact this incurs a stiff stability restriction of $O(\nu \Delta t / \Delta x^2)$ and convergence can no longer be proven.

To match what will be shown to be second order optimal spatial accuracy in SPH, we will in this work adopt the following second order incremental projection scheme (IPSPH).

$$\begin{cases} \frac{\mathbf{u}^* - \mathbf{u}^n}{\Delta t} = -\frac{1}{\rho} \nabla p^n + \frac{\nu}{2} \nabla^2 (\mathbf{u}^* + \mathbf{u}^n) & x \in \Omega, \\ \mathbf{u}^* = 0 & x \in \partial\Omega, \end{cases} \quad (9)$$

$$\begin{cases} \frac{\mathbf{u}^{n+1} - \mathbf{u}^*}{\Delta t} = -\frac{1}{\rho} \nabla (p^{n+1} - p^n) & x \in \Omega, \\ \nabla \cdot \mathbf{u}^{n+1} = 0 \\ \mathbf{u}^{n+1} \cdot \hat{n} = 0 & x \in \partial\Omega. \end{cases} \quad (10)$$

This scheme functions in an identical way: a non-divergence-free velocity is found by solving a Helmholtz problem, and the initial estimate is corrected by solving the Poisson problem in Equation 8 for $(p^{n+1} - p^n)$. By treating the viscous term with a Crank-Nicolson discretization, $O(\Delta t^2)$ convergence in velocity and $O(\Delta t)$ convergence in pressure can be proven for a Galerkin method [39]. We note that higher order convergence can be obtained by applying a non-homogeneous rotational boundary condition for the pressure. This has been investigated in an SPH context by Hosseini and Feng [34].

In an Eulerian method, the pressure at the end of the timestep is updated using the increment obtained from the Poisson problem

$$p^{n+1} = p^n + \delta p^{n+1}, \quad (11)$$

$$\delta p^{n+1} = p^{n+1} - p^n. \quad (12)$$

In a Lagrangian formulation the collocation points are not fixed between timesteps, and it is necessary to introduce the following correction to the initial pressure of the following timestep:

$$p^{n+1}(x^{n+1}) = p^n(x^n) + \delta p^{n+1}(x^n) + \nabla p^{n+1}(x^n) \cdot (\mathbf{x}^{n+1} - \mathbf{x}^n) + O((U_{max}\Delta t)^2). \quad (13)$$

2.2. Smoothed Particle Hydrodynamics: Conservation vs consistency

In meshless methods, we seek a means of constructing a continuous quasi-interpolant of a collection of data sites $\mathbf{x}_i, \mathbf{y}_i \in \mathbb{R}^n$ for $i \in 1, \dots, N$. The starting point of SPH is to approximate the desired function with a mollification [35] by taking a convolution against an approximation to the identity.

$$\langle f \rangle(\mathbf{x}) = \int_{supp(W)} f(\mathbf{y}) W(\mathbf{x} - \mathbf{y}) d\mathbf{y} \quad (14)$$

Here W is typically chosen to be a radially symmetric positive kernel with compact support and unit integral. The length scale of this support is usually characterized by the smoothing length h , which corresponds to the second moment of the function. To evaluate this integral, we seek a quadrature rule defining a set of weights V_i for each data site.

$$\tilde{f}(\mathbf{x}_i) = \sum_{j \in supp(W)} f(\mathbf{x}_j) W(\mathbf{x}_i - \mathbf{x}_j) V_j \quad (15)$$

We define these weights by applying the above quadrature rule to $\frac{1}{V}$

$$\widetilde{V^{-1}}(\mathbf{x}_i) = \sum_{j \in supp(W)} W(\mathbf{x}_i - \mathbf{x}_j) \quad (16)$$

and we have a discrete approximant that does not require the solution of an interpolation matrix (as in collocation methods) or a projection matrix (as in Galerkin methods). The basis for forming differential operators mirrors the concept of generalized derivatives in classical PDEs. The kernel plays the role of a test function and, provided the support of the kernel is contained in the domain of interest, the derivative operator can be pushed onto the kernel via integration by parts and the SPH quadrature can again be applied.

$$\langle \nabla f \rangle(\mathbf{x}) = \int_{supp(W)} f(\mathbf{y}) \nabla_{\mathbf{x}} W(\mathbf{x} - \mathbf{y}) d\mathbf{y} \quad (17)$$

$$\widetilde{\nabla f}(\mathbf{x}_i) = \sum_{j \in supp(W)} f(\mathbf{x}_j) \nabla_{\mathbf{x}_i} W(\mathbf{x}_i - \mathbf{x}_j) V_j \quad (18)$$

We note that both Equation 15 and 18 lack zeroth order consistency (i.e. neither can reproduce the exact result when applied to a constant). The nuance of SPH is how to approach this lack of consistency. To illustrate the differing approaches of different SPH schemes we present two common methods for approximating the gradient. For example, one may formally construct the gradient through the following product rule identity and apply the SPH approximation with the identity $\widetilde{fg} = \tilde{f}\tilde{g}$.

$$\nabla f = \nabla(f \cdot 1) - f \nabla 1 \quad (19)$$

$$\widetilde{\nabla f}(\mathbf{x}_i) = \sum_{j \in supp(W)} (f(\mathbf{x}_j) - f(\mathbf{x}_i)) \nabla_{\mathbf{x}_i} W(\mathbf{x}_i - \mathbf{x}_j) V_j \quad (20)$$

On the other hand, starting from the quotient rule we have the alternative expression.

$$\nabla \frac{f}{1} = \frac{1 \nabla f - f \nabla 1}{1^2} \quad (21)$$

$$\widetilde{\nabla} f(\mathbf{x}_i) = \sum_{j \in \text{supp}(W)} (f(\mathbf{x}_i) + f(\mathbf{x}_j)) \nabla_{\mathbf{x}_i} W(\mathbf{x}_i - \mathbf{x}_j) V_j \quad (22)$$

While the first approach restores zeroth order consistency, the second is anti-symmetric and gives rise to locally conservative operators when used e.g. to approximate the pressure gradient in the Euler equations. For the class of problems usually considered in SPH applications, maintaining the underlying Hamiltonian structure of the governing equations is more important and therefore the second type of approach is preferred (see [9] for a discussion of common usage). For viscous flows possessing a dissipative character, this is no longer important and the consistency properties play a fundamental role in the order of convergence of the approximation.

In the last few years, several authors have investigated the truncation error occurring in the SPH interpolation (Equation 15) and differential operators (e.g. Equations 18,20,22) [8, 18, 19, 20, 66, 67]. Generally speaking, the interpolation error can be described by the error introduced in the smoothing and in the quadrature process, i.e.

$$\|e_{interp}\| \leq \|e_{smoothing}\| + \|e_{quad}\| \quad (23)$$

while differential operators pick up an additional source of error due to particle anisotropy

$$\|e_{interp}\| \leq \|e_{smoothing}\| + \|e_{quad}\| + \|e_{anis}\| \quad (24)$$

This term vanishes when the data sites lie on a Cartesian lattice. To characterize the behavior of these sources of error, we summarize the scaling roughly as follows (see [8, 18, 19, 20] for details).

$$\|e_{smoothing}\| \sim Ch^2 \quad (25)$$

$$\|e_{quad}\| \sim C \left(\frac{h}{\Delta x} \right)^{-\beta_1} \quad (26)$$

$$\|e_{anis}\| \sim C \frac{\chi}{h^p} \left(\frac{h}{\Delta x} \right)^{-\beta_2} \quad (27)$$

where $\Delta x \sim 1/N$ denotes a lengthscale corresponding to the particle separation, β_1 and β_2 are positive integers related to the smoothness of the kernel functions, p is 1 for the gradient operator and 2 for a Laplacian operator, and χ denotes a lengthscale characterizing the departure of the particle configuration from a Cartesian lattice. To obtain a consistent and convergent result it is therefore necessary to simultaneously decrease h and $h/\Delta x$ (i.e. increase the number of neighbors interacting with each particle). For the second order operators occurring in viscous flows and in projection methods where $p = 2$, we have to increase this number of neighbors at a faster rate, giving rise to dense operators that are prohibitively expensive to invert. Additionally, if the number of neighbors per particle is fixed (as is standard practice in SPH) and N increased, then when h is decreased to the point that e_{anis} dominates, the operator diverges. To control this term a remeshing strategy is sometimes employed (see e.g. [42]); after a few timesteps the particles are interpolated back onto a Cartesian grid but the quasi-Lagrangian nature of the method is lost. Alternatively certain combinations of equations of state and variations of these divergence and gradient operators can give rise to an artificial repulsive force between particles that maintains particle ordering at the cost of an increasingly inconsistent formulation [9, 10].

In light of this analysis, we can define two goals for an SPH formulation of a high order projection method: (1) to match the second order accuracy of the smoothing error and (2) to control the quadrature and anisotropy error to obtain a convergent result.

2.3. Smoothed Particle Hydrodynamics: Local polynomial reproduction

Local polynomial reproduction, or the ability for meshless methods to exactly reproduce polynomials using compact finite shape functions, is the basis for a number of other meshfree particle methods such: as moving least squares, reproducing kernel particle method, and others [21, 22, 40, 41]. For meshless approximants and particle distributions satisfying this property, it is possible to prove [22] for an m^{th} order local polynomial reproduction $s_{f,X}$ and $f \in C^{m+1}(\Omega)$

$$|f(\mathbf{x}) - s_{f,X}(\mathbf{x})| \leq Ch^{m+1} |f|_{C^{m+1}} \quad (28)$$

where $|\cdot|_{C^{m+1}}$ denotes the H^{m+1} seminorm.

Similarly, by correcting differential operators to ensure exact results for a given degree polynomial (e.g. $\nabla_h^2 x^2 = 2$) it is possible to build consistency and higher order convergence rates back into SPH [21]. Conceptually, this plays a similar role to polynomial consistency in the finite element method. In what follows, we summarize a means of restoring linear polynomial reproduction into the gradient and Laplacian operators due to Oger [23] and Fatehi and Manzari [20] respectively. Additionally we present a means of implementing these corrections in the same asymptotic computational time as the standard SPH operators.

As a starting point for comparison we define the uncorrected 0^{th} order consistent gradient and Laplacian operators [14] as follows.

$$\nabla_0 f_i = \sum_{j \in \text{supp}(W)} (f_j - f_i) \nabla_{\mathbf{x}_i} W_{ij} V_j \quad (29)$$

$$\nabla_0^2 f_i = 2 \sum_{j \in \text{supp}(W)} \frac{f_i - f_j}{|\mathbf{r}_{ij}|} \mathbf{e}_{ij} \cdot \nabla_{\mathbf{x}_i} W_{ij} V_j \quad (30)$$

where we adopt the notation $f_i = f(\mathbf{x}_i)$, $\mathbf{r}_{ij} = \mathbf{x}_i - \mathbf{x}_j$, $\mathbf{e}_{ij} = \mathbf{r}_{ij}/|\mathbf{r}_{ij}|$ and $W_{ij} = W(\mathbf{r}_{ij})$. To achieve linear consistency we define the following corrected operators

$$\nabla_1 f_i = \sum_{j \in \text{supp}(W)} (f_j - f_i) \mathbf{G}_i \nabla_{\mathbf{x}_i} W_{ij} V_j \quad (31)$$

$$\nabla_1^2 f_i = 2 \sum_{j \in \text{supp}(W)} (\mathbf{L}_i : \mathbf{e}_{ij} \otimes \nabla_{\mathbf{x}_i} W_{ij}) \left(\frac{f_i - f_j}{|\mathbf{r}_{ij}|} - \mathbf{e}_{ij} \cdot \nabla_1 f_i \right) V_j \quad (32)$$

where $\mathbf{G}$ and $\mathbf{L}$ are both second order symmetric $n \times n$ tensors in $\mathbb{R}^n$ defined to achieve the required order consistency. We denote with r_{ij}^k , e_{ij}^k , x_i^k the k^{th} component of vectors $\mathbf{r}_{ij}$, $\mathbf{e}_{ij}$, and $\mathbf{x}_i$ respectively. Also $\partial_{i,k}$ denotes partial differentiation with respect to x_i^k . The mn component of the inverse of the gradient correction tensor $\mathbf{G}$ reads

$$(\mathbf{G}_i^{-1})^{mn} = - \sum_{j \in \text{supp}(W)} r_{ij}^m \partial_{i,n} W_{ij} V_j \quad (33)$$

To define the Laplacian correction tensor $\mathbf{L}$ we adopt the Einstein summation convention and seek a solution of the following,

$$-\delta^{mn} = \sum_{j \in \text{supp}(W)} (A_i^{kmn} e_{ij}^k + r_{ij}^m e_{ij}^n) (L_i^{op} e_{ij}^o \partial_{i,p} W_{ij} V_j) \quad (34)$$

where the third order tensor

$$A_i^{knm} = G_i^{kq} \sum_{j \in \text{supp}(W)} r_{ij}^m r_{ij}^n \partial_{i,q} W_{ij} V_j \quad (35)$$

To solve for the symmetric tensor $\mathbf{L}$ in two dimensions, we need to solve the system resulting from fixing the the pairs $(n, m) \in \{(1, 1), (1, 2), (2, 2)\}$. By precomputing for each (n, m) pair and for each particle i from Equation 35 the corresponding vector A_i^{knm} , we are left with three corresponding linear equations from Equation 34 for the coefficients of $\mathbf{L}$:

$$\begin{aligned} -1 &= \sum_j V^j (A_i^{k11} e_{ij}^k + r_{ij}^1 e_{ij}^1) [e^1 \partial_{x_1} W, (e^1 \partial_{x_2} W + e^2 \partial_{x_1} W), e^1 \partial_{x_1} W] [L_{11}, L_{12}, L_{22}]^T \\ 0 &= \sum_j V^j (A_i^{k12} e_{ij}^k + r_{ij}^1 e_{ij}^2) [e^1 \partial_{x_1} W, (e^1 \partial_{x_2} W + e^2 \partial_{x_1} W), e^1 \partial_{x_1} W] [L_{11}, L_{12}, L_{22}]^T \\ -1 &= \sum_j V^j (A_i^{k22} e_{ij}^k + r_{ij}^2 e_{ij}^2) [e^1 \partial_{x_1} W, (e^1 \partial_{x_2} W + e^2 \partial_{x_1} W), e^1 \partial_{x_1} W] [L_{11}, L_{12}, L_{22}]^T \end{aligned}$$

which can be inverted directly. In three dimensions the process is identical and results in a 6×6 matrix to invert. The total cost of precomputing A^{knm} and inverting the matrix directly is $O(3N + 3^3 N)$ in two dimensions and $O(6N + 6^3 N)$ in three dimensions and requires no additional neighbor information than the standard SPH operators; we therefore

expect for parallel implementations that this correction may increase efficiency by increasing the ratio of local to global work without increasing communication. Most importantly, we have removed the need to control e_{quad} and e_{anis} in Equation 24 with the number of neighbors per particle, and the corrected operators can achieve superior accuracy with a relatively small number of neighbors, e.g. 20 particles as opposed to 80 in two dimensions. This will give much sparser matrices with smaller bandwidth and ultimately yield more amenable matrices for solution with fast iterative solvers. We note that to date these Laplacian corrections have only been used by Fatehi et al. in application to a WCSPH formulation [56] and have not yet been applied to ISPH and the necessary Poisson solvers for achieving a divergence-free velocity field. In preliminary results we have found that for some problems, omitting these corrections can lead to the accumulation of divergence in the scheme for low to moderately resolved discretizations.

3. Spatial convergence results for representative particle distributions

Several papers investigating the uncorrected operators or their corrections generally perform their analysis by investigating the accuracy of the operators on two particle configurations: an equispaced Cartesian lattice and a perturbed lattice formed by shifting each of those particles by a random variable with distribution $U([-χ, χ] \otimes \dots \otimes [-χ, χ])$, where $χ$ denotes a scaling factor proportional to the smoothing length.

In reality this gives rise to too pessimistic of a particle distribution. In WCSPH, the equation of state serves the dual role of enforcing incompressibility and regularizing the particle distribution. In ISPH, particle regularity is enforced at the end of every timestep by shifting particles along the gradient of some indicator of particle density. Accuracy is then guaranteed by interpolating to the shifted grid using a sufficiently high order meshless interpolant. Both approaches mirror the physics that for an incompressible flow it should not be possible for mass to accumulate or voids to form.

The most commonly employed shifting strategy in ISPH originally introduced by Xu et al. [32] is to define a particle anisotropy indicator as follows.

$$\delta \mathbf{r}_i = C_{shift} \sum_{j \in supp(W)} \frac{\bar{r}_i^2}{|\mathbf{r}_{ij}|^2} \mathbf{e}_{ij} \quad (36)$$

where $C_{shift} = CU_{max}\delta t$ and $\bar{r}$ is the centroid of neighbor particle locations

$$\bar{r}_i = \frac{1}{\#\{j \in supp(W)\}} \sum_{j \in supp(W)} |\mathbf{r}_{ij}|. \quad (37)$$

To empirically characterize the operator convergence for meaningful configurations, we investigate three families of particle arrangements (Fig. 1): a Cartesian grid, a set of perturbed grids with shift scaled by $C \cdot h$ for $C \in \{0, 0.01, 0.05, 0.1, 0.5\}$, and finally a “corrected” grid generated by first shifting by $0.5 \cdot h$ and then applying the shift strategy outlined in Xu et al. [32] until the particles empirically converge to a quasi-uniform distribution representative of a typical SPH flow. Each particle configuration is used to evaluate the pointwise root mean square (RMS) error of each operator when applied to the periodic function $u(x, y) = \sin(x) \sin(y)$. This process was repeated with new random particle configurations until the mean error converged to two decimal places.

The results presented utilize a quintic spline kernel and to study convergence the ratio of the smoothing length to the initial particle spacing on the Cartesian grid is first fixed $h/\Delta x = 1.5$ (Fig. 2). In agreement with the theoretical estimates, optimal second order accuracy is achieved for the uncorrected operators on the Cartesian grid. As the configuration becomes more irregular, the gradient operator converges to an increasingly inconsistent result for the gradient operator and diverges at a rate of h^{-1} for the Laplacian. For the quasi-uniform particle configuration the error corresponds roughly to a randomly perturbed configuration of $0.03 \cdot h$. At best it is possible to achieve 5% accuracy at 24 particles per wavelength before the solution diverges, which is consistent with the $O(1\%)$ accuracy typically reported in the literature.

When the linear consistency correction is imposed, we see initial second order convergence transitioning to first order convergence with increasing anisotropy (Fig. 3) in agreement with Fatehi’s estimates. For the representative quasi-uniform configuration, the second order convergence is sustained for a system of up to one million particles with a resulting error of 10^{-4} . By repeating the experiment while changing $h/\Delta x$, we see that consistency is retained but that the transition to first order happens faster for smaller values of $h/\Delta x$ (Fig. 4). This shows a tradeoff in the

Figure 1: Laplacian approximation for representative particle distributions of Cartesian, randomly perturbed with magnitude $C = 0.1 \cdot h$, and quasi-uniform particle configurations.

Figure 2: L^2 error for uncorrected gradient (left) and Laplacian (right) operators with fixed $h/\Delta x = 1.5$.

approximation: it is possible to either have second order convergence or sparser operators. For the sparse operators and a fixed number of particles, a smaller $h/\Delta x$ implies a smaller h^2 smoothing error, and it is therefore possible for some combinations of N and $h/\Delta x$ to achieve a more accurate result with first order convergence.

When applying these corrections to a physical problem, we can therefore achieve a consistent second order result for an appropriately selected $h/\Delta x$. In practice, we expect that efficiency gains in the iterative solvers due to sparsity may outweigh the gains of second order convergence, and that this ratio will be a tunable parameter in the method.

4. Solving the singular Poisson problem

The solution of the Poisson problem with Neumann boundary data:

$$\begin{cases} -\nabla^2 u = f & x \in \Omega, \\ \partial_n u = 0 & x \in \partial\Omega, \end{cases} \quad (38)$$

Figure 3: L^2 error for corrected gradient (left) and Laplacian (right) operators with fixed $h/\Delta x = 1.5$.

is singular since the solution is unique up to the addition of an arbitrary constant. Additionally, the right hand side of the equation must satisfy the solvability constraint:

$$\int_{\Omega} f = 0.$$

Both the kernel and cokernel of the Laplacian operator contain constant functions and the system is only solvable provided f has zero mean. However, when the operators are discretized approximately as the system of equations $\mathbf{A}\mathbf{u} = \mathbf{b}$ and truncation error is introduced, the resulting system can become non-invertible. In fact, the kernel of the discretized operator A is constituted of constant vectors, but since A is in general non-symmetric, the cokernel of A is not constituted of constant vectors and the discretized right hand side $\mathbf{b}$ typically has non-zero mean. Without a guarantee that $\mathbf{b} \in \text{Im}(A)$, the resulting $\mathbf{A}\mathbf{u} = \mathbf{b}$ cannot be solved.

Following [65] we propose to solve the system by solving:

$$PAP\mathbf{u} = P\mathbf{b}, \quad \text{with } P = I - \frac{\mathbf{c}\mathbf{c}^T}{\mathbf{c}^T\mathbf{c}}$$

where $\mathbf{c}$ is any constant vector and I is the identity matrix. This projector P ensures solvability by requiring that the source term remains discretely orthogonal to the constant vector $\mathbf{c}$.

Our numerical experiments show that similar results can be achieved by solving $\mathbf{A}\mathbf{u} = P\mathbf{b}$ using a method where we replace one row i of the system with the equation $\mathbf{u}_i = 0$. After solving the system, at each time step, we take as a solution $P\mathbf{u}$ which provides a zero-mean pressure vector. We refer to these two approaches as the projection and the pinpoint methods, respectively.

In comparison between the uncorrected first order scheme and the corrected second order scheme, we note that it is particularly necessary to force the source term to have zero mean to achieve accurate results. Additionally, for long time integration with the uncorrected scheme the pressure field can decouple from the particle i where the pinpoint method is applied. This decoupling mechanism is illustrated in Fig. 5. We attribute this to the increased truncation error and lack of consistency in the uncorrected operators, further demonstrating the need for using Fatehi's consistent Laplacian discretization. This effect has not been previously observed in the literature. In general to distinguish between the two methods, the pinpoint method is much simpler to implement while the projection method is more robust. The pinpoint approach and projection one would give the same result on uniform particle set with periodic boundaries where the Laplacian matrix is symmetric, the kernel is made up of a constant, and the right hand side of the system has zero average. Due to the asymmetry of general particle arrangements however this is not the case.

Figure 4: L^2 error for corrected Laplacian operator with $h/\Delta x = 1.0$ (left) and $h/\Delta x = 2.0$ (right).

For the uncorrected scheme, the projection method was able to provide more robust results and avoid the decoupling mechanism. For the corrected scheme there was no practical difference in the results between the two approaches.

5. Serial results

5.1. Fast solvers

The SPH discretization of both the first and second order schemes require the solution of a Helmholtz equation for the predictor step and a Poisson problem to determine the pressure or pressure increment. Both of the resulting elliptic matrices are of M-form [52] for Cartesian particle distributions, for which the geometric multigrid method and its variant algebraic multigrid (AMG) have been shown to be the most efficient solvers. However, generating a hierarchy of meshes is not feasible in SPH, and the method's Lagrangian nature meshes may become severely distorted as the flow evolves. Thus, AMG methods are more suitable than geometric MG methods. For details of MG method, we refer to [50, 51] and references therein.

In our verification problems the two-dimensional serial code, we apply the AMG preconditioner implemented in the FASP linear algebra package [45]. We have investigated two different AMG methods for solving the SPH discretization including: (1) the classical AMG (C-AMG) method [48] with Gauss-Seidel smoother and V-cycle and (2) the unsmoothed aggregation AMG (UA-AMG) method [47] with Gauss-Seidel smoother and nonlinear AMLI-cycle [46]. Both methods serve as a preconditioner for a Krylov subspace iterative method (i.e. GMRES) to accelerate the overall performance of the linear solver. In Table 1 we present results of a scaling study in solving the Helmholtz problem with a large viscosity ($\nu = 10$) where the Helmholtz operator is dominated by the Laplacian term. We can observe that AMG methods are efficient for solving the Helmholtz equation, and the scaling of the computational cost is nearly optimal (i.e. the CPU time is almost linear in the degrees of freedom).

5.2. Taylor-green vortex: Temporal convergence results

To validate the second order temporal convergence of the corrected second order IPSPH scheme and compare against standard uncorrected WCSPH and ISPH, we compare results to analytical solutions for the Taylor-Green vortex. The Taylor-Green vortex has been used extensively as a validation case in meshless methods due to several of its useful properties: the flow is fully transient, the stagnation points within the flow give rise to problematic particle distributions, and the periodic boundaries remove the complications of Dirichlet boundary conditions from the results. We solve the governing equations on a periodic box of side 2π and compare to the exact solution

Figure 5: Typical decoupling of pinned particle pressure when an uncorrected scheme is used for long time integration. The particle in the bottom left corner has a pinned pressure increment of zero, but the remaining particles' pressure slowly drifts apart and decouples over time.

Table 1: Performance of AMG methods for solving Helmholtz equation in large viscosity limit with Dirichlet data

DoF	UA-AMG		C-AMG	
	#Iter.	CPU (sec.)	#Iter.	CPU (sec.)
2^{12}	7	0.03	5	0.03
2^{14}	8	0.14	6	0.16
2^{16}	8	0.64	7	0.75
2^{18}	9	2.90	8	3.42
2^{20}	9	14.15	9	16.37

$$\begin{aligned}
U_x(x, y, t) &= U_0 e^{-2\nu t} \sin(x) \cos(y) \\
U_y(x, y, t) &= -U_0 e^{-2\nu t} \cos(x) \sin(y) \\
P(x, y, t) &= \frac{U_0^2}{4} e^{-4\nu t} (\cos(2x) + \cos(2y))
\end{aligned} \tag{39}$$

To highlight the regime where viscous effects become dominant and the explicit stability restriction in WCSPH transitions from first order in time ($\Delta t = h/2U_{max}$) to second order ($\Delta t = h^2/8\nu$) [14], we investigate the regime surrounding a particle Reynolds number

$$Re_p = \frac{U_0 h}{\nu} \approx 1 \tag{40}$$

and select $U_0 = 0.1$ and $\nu = 0.1$. We note of course that for lower Reynolds numbers the viscous stability restriction will cause the explicit method to become prohibitively expensive, and we target this regime as the expected threshold where it becomes advantageous to treat the viscous term implicitly. For this highly viscous regime the particle configuration differs only slightly from the Cartesian lattice and we present results without any consistency corrections, and focus instead on the temporal convergence of the method and the effectiveness of preconditioning. The simulation is run using WCSPH for various artificial sound speeds, standard ISPH and the second order projection scheme (IPSPH) using the largest timestep allowed by their respective stability criteria until reaching $t_f = 10$. Since we are targeting a low Reynolds number regime, we therefore expect $\Delta t \sim \Delta x$, while for the explicit methods $\Delta t \sim \Delta x^2$. The pointwise RMS particle error is presented in Fig. 6(a).

The convergence results demonstrate that both the first and second order projection schemes achieve optimal accuracy, while the weakly compressible method initially achieves second order convergence due to the second order stability constraint, but convergence stagnates as the number of particles is increased. In light of Equation 4, it is clear that the artificial compressibility assumption introduces an inconsistency to the model and that convergence can only be obtained by increasing c at an extra computational cost. Numerical experiments show that the error at which this stagnation occurs scales as $\|e\|_2 \sim c^{-0.85}$.

To demonstrate the computational expense of the three methods, the total CPU time for a desired error is plotted in Fig. 6(b). From these results it is clear that while the first order ISPH method is only competitive with WCSPH within certain regimes, the second order projection method provides more efficient results uniformly, even without applying preconditioning. For the first order scheme, the computational cost scales as $t_{CPU} \sim h^{2.875}$, while the second order scheme scales as $t_{CPU} \sim h^{1.548}$. Without a preconditioner, for large N the second order projection scheme begins to scale more poorly due to the N^2 dependence of the condition number for the discrete Poisson problem. When the AMG solver is used to precondition the system, the asymptotic scaling does not show condition number dependence for the regime under investigation. The combination of high order projection schemes accelerated by fast solvers therefore shows promise as a scalable, efficient method for simulating transient viscous flow despite the overhead of inverting both the Helmholtz and Poisson problems at each timestep.

5.3. Tying everything together: A representative low Reynolds number flow

The performance of the method is evaluated under physically meaningful conditions by investigating the transient startup of flow past a periodic array of cylinders. The geometry and flow conditions are taken from the work by Morris et al. [14]. The flow geometry is presented in Fig. 7 and the parameters considered are presented in Table 2. Flow is driven past an array of cylinders of radius a by a body force F . Particles are initialized on a Cartesian grid and it is expected that a moderate number of particles will be necessary to accurately resolve the curvature of the cylinder. We validate this result using a high fidelity finite volume simulation [44] performed using 500^2 cells and a second order backward differencing scheme in time. The error is calculated by interpolating the the finite volume results at the SPH points and computing the root mean square error. In Fig. 8 the results of this process are presented for increasing resolution to demonstrate the local error distribution. We note that for well resolved simulations the error is concentrated locally near the walls in the domain.

To enforce the no-slip boundary condition, we employ the particle extrapolation technique proposed by Takeda [54]. The consistency of this approach has been discussed in the continuum limit by Macia et al. [55]. In light of their work, we note that omitting the particle extrapolation is equivalent to truncating the support of the smoothing integral and therefore introduces a local $O(h)$ error. Since at low Reynolds numbers the equations are elliptic, this error

Table 2: Periodic cylinder parameters

L (m)	$\nu (\frac{m^2}{s})$	a (m)	F ($\frac{m}{s^2}$)	Re
0.1	10^{-4}	0.02	$5 \cdot 10^{-5}$	0.01

propagates throughout the domain and gives at best first order global convergence. By applying linear extrapolation, this local error is improved to $O(h^2)$ and matches the smoothing and time-stepping error.

For the current problem, we initialize the particle distribution on a Cartesian lattice and denote ghost particles as those falling within the cylinder. The mirroring distances are computed in an identical manner to those in [14]. We note that Hashemi has proposed an alternative boundary condition treatment that does not require ghost particles [56]. We retain the reflecting ghost particle technique for the current work since it is used more commonly in the SPH literature and is consistent with the differential operator corrections and gives second order spatial accuracy. Results are compared at a final time of $t = 20$. Again we use the largest stable timestep allowed by each method, meaning that for the implicit methods $\Delta t \sim \Delta x$, while for the explicit methods $\Delta t \sim \Delta x^2$.

Results are presented for: the standard WCSPH approach for $c = 5U_{max}$ and $c = 20U_{max}$, for the first order ISPH approach, and for the second order incremental pressure correction scheme using $h/\Delta x = 1.5$. The gradient and Laplacian corrections were then applied to the second order scheme to demonstrate their effect on computational efficiency. In Fig. 9 results from the second order projection method are compared to the finite volume results demonstrating excellent agreement. The error analysis presented in Fig. 10(a) shows that both projection approaches achieve their formal orders of accuracy after the cylinder is fully resolved. For resolution larger than 48^2 particles the convergence of the second order method stagnates as the method has converged to the same level of resolution as the finite volume method.

The WCSPH results on the other hand converge at roughly first order for both $c = 5U_{max}$ and $c = 20U_{max}$. From these results it is clear that with only minimal resolution (32^2 particles) the second order scheme achieves more accurate answers than either of the standard SPH schemes do with full refinement (192^2 particles). Further, Fig. 10(b) shows that while WCSPH outperforms ISPH despite its stiff timestep restriction, the second order projection method outperforms both schemes by several orders of magnitude for a fixed error.

When the particle corrections are enforced with the same smoothing length, Fig. 10(b) demonstrates that negligible cost is added to the method. When the ratio is reduced to $h/\Delta x = 1.0$, an approximately four times further speed-up is achieved at the expense of reverting to first order spatial convergence, in agreement with the results of Section 3. We therefore see that by combining fast solvers, an efficient scheme and polynomial reproduction we can achieve a better answer in less than a second than what is possible in hours using traditional SPH approaches.

Having demonstrated the optimal convergence properties and practical computational advantages of the scheme, we demonstrate that the scheme remains robust and stable for long time integration. For flows evolved over long integration times or for higher Reynolds number flows, the particles become disordered and it is necessary to apply the correction procedure described in Equation 36 [32]. In these regimes the requirements for spatial accuracy dominates the temporal accuracy and it can be assumed asymptotically that the large timesteps taken using the projection scheme will not affect the error. However, while the standard ISPH scheme employs a backward Euler integration scheme, Crank-Nicolson is a symplectic time integrator and therefore introduces no numerical dissipation that might stabilize the method. To demonstrate the stability properties of the scheme, the periodic cylinder case is evolved forward in time to two full traversals of the domain by the fastest particle at $t = 500$. Fig. 11 shows that the particles are fully disordered, although the particle shifting strategy ensures a quasi-isotropic configuration. The approach to steady state of the maximum fluid velocity in Fig. 12 heuristically demonstrates the stability of the method. For $N = 48^2$ particles similar results were confirmed up to $t = 5000$.

6. Parallel Implementation

We have developed a parallel 3D implicit SPH code using existing libraries to leverage the proposed second order projection scheme for large scale simulations. LAMMPS is originally developed for molecular dynamics simulations

and provides computational infrastructures for massive parallel particle simulations [58]. However, LAMMPS does not have the capability for implicit time integration nor the parallel linear algebra which is required for solving Poisson and Helmholtz problems in our formulation. For an efficient parallel solutions of those linear problems, we have developed an interface integrating LAMMPS with packages from the Trilinos library [57]. Specifically, we use the Epetra package [63] for distributed memory linear algebra, the Belos package [60] for a GMRES linear solver, and the ML package [62] for smoothed aggregation AMG preconditioners for block GMRES. The interfaced code lets LAMMPS handle the particle simulation data and Trilinos handle the distributed memory linear solves, allowing each code to do what it was designed to do well. As both LAMMPS and the Trilinos libraries have demonstrated massively parallel scalability, our interfaced code inherits that performance. To the best of the authors' knowledge, the computational results presented below represent the largest implicit SPH simulations ever performed. As such, this parallel implementation of implicit SPH represents a fundamentally new capability for the investigation of relevant physical phenomena with a level of fidelity never before possible.

For both the Helmholtz and Poisson systems, we use ML's smoothed aggregation algebraic multigrid with 4 sweeps (pre and post) of Gauss-Seidel as a smoother on each level except the coarsest one, on which a serial direct solve is performed. To improve parallel performance we repartition levels onto a subset of processors using Zoltan when the number of unknowns per core falls under 500 [64]. In each case ML is used to precondition GMRES (Poisson) or block GMRES (Helmholtz) where the convergence tolerance on the Krylov method is set to 10^{-10} .

We first replicate the convergence study described in Section 5, but extrude the problem domain to be fully 2π periodic in the z dimension, so that the computation is fully 3D. The corresponding analytic solution is constant in z and given by Equation 39. Simulations are run on NERSC's Cray XE6 (Hopper). First, results corresponding to $h = 1.5\Delta x$ are presented to investigate scaling for the regime where second-order convergence was previously observed in the Laplacian operators (Section 3). As expected, second order convergence was observed, as reflected in Fig. 13. Parallel scalability for up to 134 million particles computed using 32,768 cores is reported in Fig. 14. The number of iterations per linear solve is shown on the left of Fig. 14. It shows that the iteration counts grow modestly with increasing problem size for both the Helmholtz and Poisson solves. The right side of Fig. 14 shows the weak scaling efficiency per linear solver iteration, namely the cost per iteration at n cores divided by the cost per iteration at 64 cores. The ideal method would have a weak scaling efficiency of one, but in practice we expect that to decrease slightly as the problem size increases. While AMG theoretically converges in a constant number of iterations, in practice the number of iterations required for convergence increases mildly for massive scale problems. The efficiencies shown for the Helmholtz solver are very good, as they are near and sometimes even above one. Efficiencies greater than one are usually caused by machine jitter or network effects and should not be taken to imply that the underlying algorithms are more efficient on larger core counts. The Poisson solver efficiencies are more typical, declining to about 0.7 at 32,768 cores.

As previously noted in Section 3, the consistency operators allow for a much smaller support radius than that typically used in the SPH literature. By setting $h = 0.8\Delta x$, the increased sparsity of the operators leads to better memory bandwidth utilization and we see a four-fold decrease in the total wall-time of the method. As discussed previously, this gives more accurate results since the truncation error scales with $O(h^2)$ and transitions to $O(h)$ for large numbers of particles. This leads to the somewhat counter-intuitive conclusion that the extra linear algebra necessary in inverting the small local matrices for the consistency corrections actually leads to a much faster method, since the global matrices become cheaper to solve. The scaling results corresponding to this reduced support radius are presented in Fig. 15. While the number of iterations is similar to what is shown in the previous result, the parallel efficiency per iteration illustrated in the right of Fig. 15 shows bigger fluctuations. This is likely because the reduced time complexity due to the sparser operators amplifies the network or hardware noise.

Next, to highlight the capability of the code to simulate complex 3D geometries, we consider flow past a 3D periodic array of spheres, similar to the 2D example presented in Section 5.3. The spheres are arranged in a body-centered cubic structure within a periodic box, as shown in Fig. 16. The cube edge length L , the spheres' radius a , the viscosity ν and the forcing term F are the same as the 2D example previously given in Table 2. To enforce the Takeda/Morris no-slip boundary condition requires the evaluation of the perpendicular distances to walls for fluid particles interacting with ghost particles. In increasingly complex boundary configurations this becomes cumbersome and in general requires some sort of piecewise spline representation of the boundaries. To circumvent this we have implemented the smoothed perpendicular distance approximation introduced by Holmes et al. [43]. This technique approximates the perpendicular distances using only information corresponding to the distribution and ID (either

interior or ghost particle) of nearest neighbors and preserves a truly meshfree boundary representation. To validate these results we present a comparison with a finite volume simulation by analogy with Section 5.3. Grid convergence for the benchmark was conducted by performing a refinement study for 9408, 71632, and 768000 hexahedra cell meshes. In Fig. 17 and Fig.18 we compare using results from an SPH simulation with $N = 64^3$ particles with smoothing length $h = 0.8 * \Delta x$ at time $t = 20s$. To highlight the agreement between the results, the Shepard interpolant

$$u_{shep}(\vec{x}) = \frac{\sum_j u_j (1 - \frac{\|\vec{x} - \vec{x}_j\|}{h_{shep}})}{\sum_k (1 - \frac{\|\vec{x} - \vec{x}_k\|}{h_{shep}})}$$

with support $h_{shep} = 0.0025$ is used to compare SPH streamwise velocity and pressure to the finite volume results along the line starting at $y = 0.05, z = 0.05$ in the streamwise direction. Fig. 19 demonstrates excellent agreement between the two results, with a small discrepancy near the center particle that is attributable to the truncation of the support of the Shepard interpolant.

7. Conclusions and future work

The ISPH interpretation of smoothed particle hydrodynamics provides a Lagrangian framework and therefore a powerful modeling tool for studying a variety of fluid mechanics problems. When using SPH to study transient viscous flow however, several inconsistencies arise if the standard methodology is applied and must be addressed to achieve a consistent approach. While deficiencies in the accuracy and consistency of the operators may be safely ignored in favor of conservation properties for nearly inviscid problems, the dissipative nature of low Reynolds number flow necessitates the construction of consistent differential operators to achieve an accurate answer. If the equations of motion are solved using an approximate splitting scheme, the inconsistency in the standard Laplace operator will lead to large errors in the divergence field for problems requiring a large amount of resolution. If the artificial compressibility assumption is applied, an additional inconsistency appears as a modelling error. We have demonstrated in this work that in the process of correcting these problems and restoring consistency to the discrete model, the resulting approximation is orders of magnitude more efficient than the classical approach. By combining an accurate scheme, a fast solver, a particle isotropy correction and local polynomial reproduction we attain a different interpretation of SPH that is better suited to low Reynolds number problems. Philosophically, the approximation resembles higher order meshless methods (such as moving least squares, reproducing kernel particle method, meshfree Galerkin, etc.) but maintains the efficiency of the collocation-like nature of SPH. While these other meshless methods typically need to incorporate some type of mesh to evaluate Galerkin forms of the meshless approximants, the current approach maintains the efficiency and truly meshfree nature of SPH.

The efficiency of this method relies on preserving a high enough order of convergence that large timesteps can be used to offset the computational overhead of the linear solvers without introducing additional error. With this in mind, the standard techniques used in the literature need to be confirmed or altered so as to not contaminate the second order convergence. For example many of the approaches to applying boundary conditions or studying multiphase flow apply physics-based justifications or are derived with the end goal of preserving conservation rather than consistency. Non-periodic boundary conditions will require a consistent introduction of particles at inlets, and the particle splitting/conglomeration strategies used to apply adaptivity to SPH need to be done using second order interpolation.

In the current work for simplicity we have neglected to investigate the effect of applying a non-homogeneous rotational pressure boundary condition instead of the simpler Neumann boundary condition used here. This is widely recognized as necessary for avoiding the introduction of numerical boundary layers in the divergence due to the splitting [36], and has been investigated by Hosseini et al. in an ISPH scheme but without using a consistent splitting operator [34]. While this is certainly a necessary future work, it is interesting to point out that the current scheme achieves good agreement with finite volume results, suggesting that perhaps a consistent Poisson solver is more important in accurately evaluating drag forces than the rotational boundary condition.

There is a large class of problems where classical WCSPH, and particularly recent reformulations addressing inconsistencies in the artificial compressibility assumption [10, 68], are particularly suited to achieve accurate and robust results. In particular, for inviscid free surface problems these explicit methods have demonstrated excellent

results [67] and it is not clear whether the approach outlined in this work can be extended to robustly handle these problems. For example, in free surface flows when pinch-off events are under-resolved, there are no guarantees that the correction matrices be invertible. There are alternative multiphase methods in which the interface is not explicitly tracked but instead treated diffusely [69], but it remains to be seen whether the corrections used in the current paper will remain stable in the presence of large density ratios or at higher Reynolds numbers. It would also be interesting to see how the accuracy and efficiency of the current approach would compare to these more modern WCSPH formulations.

In addition to achieving more accurate and faster results, a key feature of the current approach is its scalability. With a consistent approximation of the Navier-Stokes equations it becomes possible to achieve arbitrary accuracy for a given problem by simply increasing resolution. For a large class of problems (i.e. turbulent flows, fluid structure interaction, non-linear wave propagation, etc.) a direct numerical simulation requires a high level of fidelity and the ability to achieve accuracy over a range of length scales while maintaining scalability. To this end, the current work demonstrates that the efficiency gains achieved by the combination of schemes and corrections above translates to a fast massive-scale parallel implementation while maintaining scaling up to large numbers of particles. We expect that moving into the future, the approach and software developed in this paper will provide a robust and powerful foundation for studying multiphysics problem that cannot be approached with current SPH implementations.

Acknowledgments

We acknowledge the extremely helpful contributions from George Karniadakis at Brown University in discussing projection schemes, Chris Siefert at Sandia National Laboratory in setting up the algebraic multigrid solvers in Trilinos, and Wenxiao Pan at Pacific Northwest National Laboratory for helpful discussions regarding SPH. This material is based upon work supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program as part of the Collaboratory on Mathematics for Mesoscopic Modeling of Materials (CM4), under Award Number DE-SC0009247. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the U.S. Department of Energy under contract DE-AC04-94-AL85000.

References

- [1] Monaghan, J. J. "Smoothed particle hydrodynamics." *Reports on progress in physics* 68.8 (2005): 1703.
- [2] Monaghan, J. J. "Smoothed particle hydrodynamics and its diverse applications." *Annual Review of Fluid Mechanics* 44 (2012): 323-346.
- [3] Monaghan, J. J., and R. A. Gingold. "Shock simulation by the particle method SPH." *Journal of Computational Physics* 52.2 (1983): 374-389.
- [4] Gingold, R. A., and J. J. Monaghan. "Kernel estimates as a basis for general particle methods in hydrodynamics." *Journal of Computational Physics* 46.3 (1982): 429-453.
- [5] Gingold, R.A. and Monaghan, J.J. "Smoothed particle hydrodynamics: theory and application to non-spherical stars." *Monthly Notices Royal Astronomical Society*, Vol. 181, pp. 375-389, (1977).
- [6] Lucy, L.B. "A numerical approach to the testing of fusion process." *Astronomical Journal*, Vol. 88, pp. 1013-1024, 1977.
- [7] Bonet J, Lok T-S. "Variational and momentum preserving aspects of smoothed particle hydrodynamics formulations. *Comput. Methods Appl. Mech. Eng.* 180:9711, (1999).
- [8] Quinlan, Nathan J., Mihai Basa, and Martin Lastiwka. "Truncation error in meshfree particle methods." *International Journal for Numerical Methods in Engineering* 66.13 (2006): 2064-2085.
- [9] Price, Daniel J. "Smoothed particle hydrodynamics and magnetohydrodynamics." *Journal of Computational Physics* 231.3 (2012): 759-794.
- [10] Adami, S., X. Y. Hu, and N. A. Adams. "A transport-velocity formulation for smoothed particle hydrodynamics." *Journal of Computational Physics* 241 (2013): 292-307.
- [11] Colagrossi, Andrea, and Maurizio Landrini. "Numerical simulation of interfacial flows by smoothed particle hydrodynamics." *Journal of Computational Physics* 191.2 (2003): 448-475.
- [12] Morris, Joseph P. "Simulating surface tension with smoothed particle hydrodynamics." *International Journal for Numerical Methods in Fluids* 33.3 (2000): 333-353.
- [13] Grenier, N., Antuono, M., Colagrossi, A., Le Touze, D., and Alessandrini, B. (2009). "An Hamiltonian interface SPH formulation for multi-fluid and free surface flows." *Journal of Computational Physics*, 228(22), 8380-8393.
- [14] Morris, Joseph P., Patrick J. Fox, and Yi Zhu. "Modeling low Reynolds number incompressible flows using SPH." *Journal of Computational Physics* 136.1 (1997): 214-226.
- [15] Hu, X. Y., and N. A. Adams. "A multi-phase SPH method for macroscopic and mesoscopic flows." *Journal of Computational Physics* 213.2 (2006): 844-861.
- [16] Filipovic, Nenad, Milos Ivanovic, and Milos Kojic. "A comparative numerical study between dissipative particle dynamics and smoothed particle hydrodynamics when applied to simple unsteady flows in microfluidics." *Microfluidics and nanofluidics* 7.2 (2009): 227-235.

- [17] Mittal, Rajat, and Gianluca Iaccarino. "Immersed boundary methods." *Annu. Rev. Fluid Mech.* 37 (2005): 239-261.
- [18] Graham, David I., and Jason P. Hughes. "Accuracy of SPH viscous flow models." *International journal for numerical methods in fluids* 56.8 (2008): 1261-1269.
- [19] Basa, Mihai, Nathan J. Quinlan, and Martin Lastiwka. "Robustness and accuracy of SPH formulations for viscous flow." *International journal for numerical methods in fluids* 60.10 (2009): 1127-1148.
- [20] Fatehi, R., and M. T. Manzari. "Error estimation in smoothed particle hydrodynamics and a new scheme for second derivatives." *Computers and Mathematics with Applications* 61.2 (2011): 482-498.
- [21] Belytschko, T., et al. "On the completeness of meshfree particle methods." *International Journal for Numerical Methods in Engineering* 43.5 (1998): 785-819.
- [22] Wendland, Holger. *Scattered data approximation*. Vol. 17. Cambridge: Cambridge University Press, 2005.
- [23] Oger, G., M. Doring, B. Alessandrini, and P. Ferrant. "An improved SPH method: Towards higher order convergence." *Journal of Computational Physics* 225.2 (2007): 1472-1492.
- [24] Chorin, Alexandre Joel. "A numerical method for solving incompressible viscous flow problems." *Journal of Computational Physics* 2.1 (1967): 12-26.
- [25] Lastiwka, Martin, Mihai Basa, and Nathan J. Quinlan. "Permeable and nonreflecting boundary conditions in SPH." *International journal for numerical methods in fluids* 61.7 (2009): 709-724.
- [26] Dominguez, J. M., Crespo, A.J., and Gomez-Gesteira, M. "Optimization strategies for CPU and GPU implementations of a smoothed particle hydrodynamics method." *Computer Physics Communications* 184.3 (2013): 617-627.
- [27] Dominguez, J. M., Crespo, A.J., Valdez-Balderas, D., Rogers, B. D., and Gomez-Gesteira, M. (2013). "New multi-GPU implementation for smoothed particle hydrodynamics on heterogeneous clusters". *Computer Physics Communications*, 184(8), 1848-1860.
- [28] Cummins, Sharen J., and Murray Rudman. "An SPH projection method." *Journal of computational physics* 152.2 (1999): 584-607.
- [29] Braess, Dietrich. *Finite elements: Theory, fast solvers, and applications in solid mechanics*. Cambridge University Press, 2001.
- [30] Hu, X. Y., and N. A. Adams. "An incompressible multi-phase SPH method." *Journal of Computational Physics* 227.1 (2007): 264-278.
- [31] Lee, E. S., Moulinec, C., Xu, R., Violeau, D., Laurence, D., and Stansby, P. (2008). "Comparisons of weakly compressible and truly incompressible algorithms for the SPH mesh free particle method." *Journal of Computational Physics*, 227(18), 8417-8436.
- [32] Xu, Rui, Peter Stansby, and Dominique Laurence. "Accuracy and stability in incompressible SPH (ISPH) based on the projection method and a new approach." *Journal of Computational Physics* 228.18 (2009): 6703-6725.
- [33] Dukowicz, John K., and Arkady S. Dvinsky. "Approximate factorization as a high order splitting for the implicit incompressible flow equations." *Journal of Computational Physics* 102.2 (1992): 336-347.
- [34] Hosseini, S. Majid, and James J. Feng. "Pressure boundary conditions for computing incompressible flows with SPH." *Journal of Computational Physics* 230.19 (2011): 7473-7487.
- [35] Evans, Lawrence C. "Partial differential equations." (1998).
- [36] Karniadakis, George Em, Moshe Israeli, and Steven A. Orszag. "High-order splitting methods for the incompressible Navier-Stokes equations." *Journal of computational physics* 97.2 (1991): 414-443.
- [37] Guermond, J. L., Peter Mineev, and Jie Shen. "An overview of projection methods for incompressible flows." *Computer methods in applied mechanics and engineering* 195.44 (2006): 6011-6045.
- [38] A.J. Chorin, Numerical solution of the Navier-Stokes equations, *J. Math. Comp.* 22 (1968) 745-762.
- [39] J. Shen. On error estimates of projection methods for the Navier-Stokes equations: second-order schemes. *Math. Comput.* 65 (215) (1996)
- [40] S. Li, W. K. Liu. *Meshfree Particle Methods*. Springer, 2004.
- [41] Fasshauer, Gregory E. *Meshfree approximation methods with MATLAB*. Vol. 6. World Scientific, (2007).
- [42] Chaniotis, A. K., D. Poulidakos, and P. Koumoutsakos. Remeshed smoothed particle hydrodynamics for the simulation of viscous and heat conducting flows. *Journal of Computational Physics* 182.1 (2002): 67-90.
- [43] Holmes, David W., John R. Williams, and Peter Tilke. "Smooth particle hydrodynamics simulations of low Reynolds number flows through porous media." *International Journal for Numerical and Analytical Methods in Geomechanics* 35.4 (2011): 419-437.
- [44] Weller, H. G., Tabor, G., Jasak, H., and Fureby, C., A Tensorial Approach to Computational Continuum Mechanics Using Object-Oriented Techniques, *Computers in Physics*, Vol. 12, No. 6, (1998), pp. 620-63.
- [45] FASP Package. <http://fasp.sourceforge.net/>.
- [46] X. Hu, P. Vassilevski, and J. Xu. "Comparative convergence analysis of nonlinear multi-cycle multigrid." *SIAM Journal on Numerical Analysis*, 51(2):1349-1369, (2013).
- [47] H. Kim, J. Xu, and L. Zikatanov. "A multigrid method based on graph matching for convection-diffusion equations." *Numer. Linear Algebra Appl.*, 10(1-2):181-195, (2003).
- [48] J. Ruge and K. Stüben. "Algebraic multigrid." In *Multigrid methods*, volume 3 of *Frontiers Appl. Math.*, pages 73-130. SIAM, Philadelphia, PA, (1987).
- [49] K. Stüben. "An introduction to algebraic multigrid." In U. Trottenberg, C. Oosterlee, and A. Schüller, editors, *Multigrid*, pages 413-528. Academic Press, San Diego, CA, (2001).
- [50] U. Trottenberg, C. Oosterlee, and A. Schüller. *Multigrid*. Academic Press, London, (2001).
- [51] P. Vassilevski. *Multilevel Block Factorization Preconditioners: Matrix-based Analysis and Algorithms for Solving Finite Element Equations*. Springer Verlag, (2008).
- [52] Saad, Yousef. *Iterative methods for sparse linear systems*. Siam, (2003).
- [53] J. Xu. "Iterative methods by space decomposition and subspace correction." *SIAM Rev.*, 34(4):581-613, (1992).
- [54] Takeda, Hidenori, Shoken M. Miyama, and Minoru Sekiya. "Numerical simulation of viscous flow by smoothed particle hydrodynamics." *Progress of Theoretical Physics* 92.5 (1994): 939-960.
- [55] F. Macia, M. Antuono, L.M. Gonzalez, A. Colagrossi. "Theoretical analysis of the no-slip boundary condition enforcement in SPH methods." *Progress of Theoretical Physics* 125(6). 1091-1121
- [56] Hashemi, M. R., R. Fatehi, and M. T. Manzari. "A modified SPH method for simulating motion of rigid bodies in Newtonian fluid flows."

- International Journal of Non-Linear Mechanics* 47.6 (2012): 626-638.
- [57] M. Heroux, R. Bartlett, V. Howle, R. Hoekstra, J. Hu T. Kolda, R. Lehoucq, K. Long, R. Pawlowski, E. Phipps, A. Salinger, H. Thornquist, R. Tuminaro, J. Willenbring and A. Williams. *An Overview of Trilinos*. Sandia National Laboratories, SAND2003-2927, (2003).
 - [58] S. Plimpton. *Fast Parallel Algorithms for Short-Range Molecular Dynamics*. J Comp Phys, 117, 1-19, (1995).
 - [59] G.C. Ganzenmüller, M.O. Steinhauser, P.V. Liedekerke. *A guide to the SPH-USER package*. (2011).
 - [60] E. Bavier, M. Hoemmen, S. Rajamanickam, H. Thornquist. *Amesos2 and Belos: Direct and iterative solvers for large sparse linear systems*. Scientific Programming, (2012), Volume 20, Issue 3
 - [61] Randles, P. W., and L. D. Libersky. "Smoothed particle hydrodynamics: some recent improvements and applications." *Computer methods in applied mechanics and engineering* 139.1 (1996): 375-408.
 - [62] M.W. Gee and C.M. Siefert and J.J. Hu and R.S. Tuminaro and M.G. Sala. *ML 5.0 Smoothed Aggregation User's Guide*. Sandia National Laboratories, SAND2006-2649, (2006)
 - [63] M.A. Heroux. *Epetra Performance Optimization Guide*. Sandia National Laboratories, SAND2005-1668, (2009)
 - [64] K. Devine and E. Boman and R. Heaphy and B. Hendrickson and C. Vaughan. *Zoltan Data Management Services for Parallel Dynamic Applications*. Computing in Science and Engineering 4(2): 90-97, (2002).
 - [65] P. Bochev and R. B. Lehoucq. *On finite element solution of the pure Neumann problem*. SIAM Rev, 47, (2001).
 - [66] M. Antuono, B. Bouscasse, A. Colagrossi, S. Marrone, A measure of spatial disorder in particle methods, *Computer Physics Communications*, 185, 2609-2621 (2014).
 - [67] D. Le Touze, A. Colagrossi, G. Colicchio, M. Greco, A critical investigation of Smoothed Particle Hydrodynamics applied to problems with free surfaces, *International Journal of Numerical Methods in Fluids*, 73, 660-691, (2013).
 - [68] Marrone, S., Colagrossi, A., Antuono, M., Colicchio, G., and Graziani, G. (2013). "An accurate SPH modeling of viscous flows around bodies at low and moderate Reynolds numbers". *Journal of Computational Physics*, 245, 456-475.
 - [69] Tartakovsky, Alexandre, and Paul Meakin. "Modeling of surface tension and contact angles with smoothed particle hydrodynamics." *Physical Review E* 72.2 (2005): 026301.

(a)

(b)

Figure 6: Convergence results and computational cost for Taylor-Green vortex case

Figure 7: Geometry for periodic cylinder case

(a)

(b)

(c)

Figure 8: Local error distribution in the velocity magnitude for the periodic cylinder case with increasing resolution; number of particles (a) 16^2 , (b) 32^2 , (c) 64^2 . Error is calculated by linearly interpolating finite volume result to SPH particle locations and calculating absolute value.

Figure 9: Comparison of finite volume (top) and second order projection SPH (bottom) results for streamwise velocity component u_x and pressure p for 48^2 particles with $h = 1.5 \cdot \Delta x$.

Figure 10: (a) Convergence results and (b) computational cost per RMS error for periodic cylinder case

Figure 11: Particle configuration for long time integration at $t = 500$

Figure 12: Maximum velocity for long time integration

Figure 13: Taylor Green Vortex 3D test: velocity RMS error as a function of N^3 total particles .

Figure 14: Weak scalability results for Helmholtz (blue) and Poisson (red) solves with a smoothing length $h = 1.5\Delta x$, having 4096 particle per core. Left: number of linear iterations per solve as a function of number of cores. Right: weak scaling efficiency per iteration as a function of number of cores.

Figure 15: Weak scalability results for Helmholtz (blue) and Poisson (red) solves with a smoothing length $h = 0.8\Delta x$, having 4096 particle per core. Left: number of linear iterations per solve as a function of number of cores. Right: weak scaling efficiency per iteration as a function of number of cores.

Figure 16: Streamlines starting from selected points colored with magnitude of velocity generated from SPH results.

Figure 17: Streamwise velocity (m/s) and pressure (Pa) in plane $z = 0.05$ with comparison to finite volume results. Mesh shown for 71632 cell case to demonstrate mesh structure.

Figure 18: Normal velocity (m/s) and pressure (Pa) in plane $z = 0.0$ with comparison to finite volume results.

Figure 19: Streamwise probe along center of domain comparing interpolated SPH and finite volume velocity and pressures normalized by maximum value.