Exploring multi-resolution particle CFD methods

by James Crist
EnginSoft USA

Particle-based CFD (computational fluid dynamics) methods have become very popular in recent years due to the simplicity of the model configuration process and the ability to solve free surface problems such as splashing. One of the methods for free-surface problems is called moving particle simulation (MPS) and is frequently used in the automotive industry to evaluate splash patterns and churning losses in gearboxes.

The method assumes incompressibility (constant density) of the fluid which allows it to use larger time steps than most other particle methods. Historically the disadvantage of particle methods is related to their ease of use: the end user is able to preprocess the model by specifying only a fixed particle size that the software uses to discretize the model, thus providing a constant spatial resolution of the domain.

This particle size must be set according to the smallest feature, so large models with small features require a large number of particles. As the number of particles increases, additional memory and computing power are required which can make problems impractical to solve. Several approaches have been developed to manage this problem, each with its own advantages and disadvantages.

The fundamental challenge is that particle methods can guarantee conservation of mass and momentum by virtue of resolving fluid with discrete, persistent particles, but this guarantee must be relaxed to implement multiple particle sizes in a single model.

The MPS approach allows one or more regions to be defined in which the particle size is a fraction of the global size. Around these regions a transition zone is created in which the particles can split (or merge) so that mass and momentum are conserved.

This paper will discuss the challenges inherent in variable-resolution particle methods, some existing methods, and the new approach taken by MPS.

History

Particle CFD methods were developed to solve a class of problems difficult for traditional mesh-based codes. Smooth particle hydrodynamics (SPH) was initially developed for the astrophysics community to analyse the formation of astronomical objects such as stars or solar systems. In these models, density varies from a vacuum to the interior of a star, so SPH particles are constructed to manage large density gradients. An SPH particle is a distribution of mass centred around a point. To solve for the system's mass distribution, one simply takes a superposition of all particle mass distributions [1].

MPS was introduced by Kashizuka and Oka as a faster particle method for purely incompressible liquid flows [2]. It describes a particle as a sphere of fluid where particles are not allowed to overlap. This approach has found success in modelling free-surface problems such as gearbox oil splashing where meshing is a time-consuming process.

In both cases these methods discretize a problem into an array of fixed-size particles and then allow the particles to move based on their interactions with their neighbours. Moving boundaries are easily represented by special wall nodes which interact with fluid particles. Motion can be applied to these wall nodes without the need for complex moving or deformable meshes since a particle is not required to exist on the surface of the wall.

Conservation of energy

Unlike the Eulerian approach which has fixed connections between nodes over many time steps, particle CFD methods dynamically evaluate the interactions between a particle and its neighbours at each time step. At each time step a particle will identify its neighbours as particles within a certain distance, or within a local sphere of influence. This sphere can be described as a ratio of particle diameter which introduces the first challenge when implementing a particle solver with multiple particle diameters: the sphere will change with particle size.

If one chooses to make the sphere of influence constant, relative to particle size, then smaller particles will have a smaller sphere of influence. Close to the boundary of a high-resolution region there will be pairs of large and small particles where the small particle is in the sphere of influence of the large particle but not vice versa, leading to unbalanced forces and ultimately to the creation or removal of energy.

If one chooses to make the sphere of influence constant, regardless of particle size, then smaller particles will have more neighbours than larger particles. As the number of neighbours increases, more calculations are required to evaluate the net force on a given particle, increasing the computational cost. In this case a multi-resolution model can have more total neighbours than the same model with entirely smaller particles if the ratio of small-to-large particles is sufficiently high.

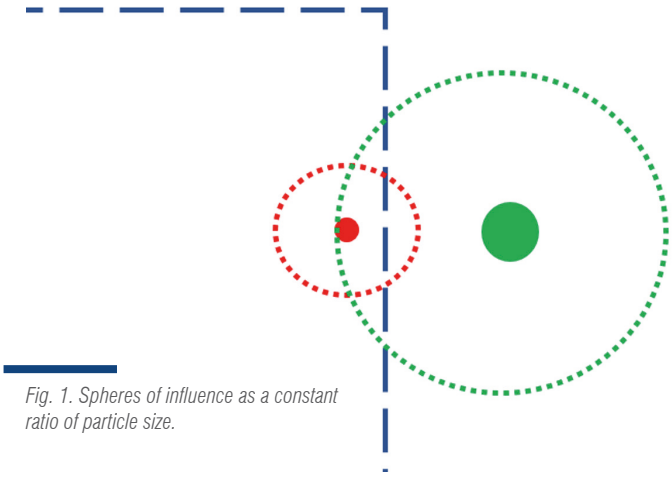

Fig. 1. Spheres of influence as a constant ratio of particle size.

The MPS multi-resolution implementation adopts a hybrid approach by introducing a transition region surrounding the region of smaller particles. Within the transition region the small and large particle problems are separated, with the properties of the large particles acting as boundary conditions for the small-particle problem. The properties of small particles are smoothed and applied in a similar way to large particles [3].

Conservation of mass

The next challenge that arises when implementing a multi-resolution method is conservation of mass. When defining a multi-resolution model it is necessary to specify certain regions that can contain smaller particles. As the simulation evolves particles will cross the boundary of this region and must be converted from small to large or vice versa.

Clearly this cannot be done one for one particle while conserving mass. Instead, some fixed ratio of large and small particles are exchanged. Even if this ratio is an integer number of small particles per single large particle it may not be possible to conserve mass. A problem arises when a single particle, or an excess particle, leaves the high-resolution region in a given time step the solver must relax mass conservation or the high-resolution boundary.

The MPS implementation chooses to do the latter by allowing small particles to exist in the transition region described above. In the time it takes for a particle to travel the length of the transition region, it becomes more likely to have enough neighbours leaving the region to constitute a large particle [3].

Conservation of momentum

The case of a single particle leaving a high-resolution region can be extended to highlight another challenge with variable-resolution particle methods. When a group of small particles is replaced by a large particle, multiple momentum vectors must

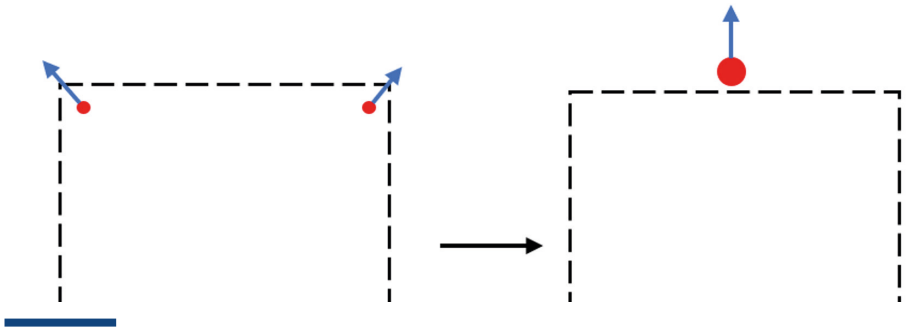

Fig. 2. Merging particles with opposing velocity vectors.

be replaced by a single vector. For a group of neighbouring particles, they are likely to have similar velocity vectors but this will not always be the case. If a small number of particles leave the high-resolution region, the resulting velocity vector may lead to non-physical behaviour.

The MPS implementation limits this behaviour by interpolating the velocity and other properties of the deleted particles to a neighbouring point representing a large particle. This interpolation is distance-weighted so a relatively distant particle will have little impact on the direction of travel of the larger particle.

Time steps

Both SPH and MPS have stability criteria for time steps based on particle size, with a typical rule for MPS being

$$\Delta t = C\frac{l_0}{u_{max}}$$

where C is the Courant number, l_0 is the particle diameter, and u_{max} is the maximum velocity in the system. It is clear from this equation that small particles require a proportionally smaller time step than large particles.

Lowering the time step of all particles based on the minimum size would result in unnecessary calculations so the MPS implementation introduces additional time steps.

The time step for the smaller particles is reduced by the same factor as the particle size resulting in multiple sub-steps of small particles between each full-time step of large and small particles. Since the sub-steps can result in unstable configurations of overlapping particles in the transition region an additional step is added. Particle displacement is applied to the overlapping particles to move them slightly into a more stable configuration [3].

Conclusion

While particle methods have many advantages over their mesh-based counterparts their strengths can become weaknesses. The trade-offs that allow them to efficiently simulate high density gradients and sparsely populated domains make changing the resolution across the domain challenging.

The MPS approach involves decoupling large and small particle regions with a transition zone between them. Within the small-particle regions the time step is reduced locally to preserve stability without slowing down the rest of the model.

For more information:
James Crist - EnginSoft USA
james@enginsoftusa.com

References

[1] EnginSoft USA, "Smoothed Particle Hydrodynamics" [Online]. Available: enginsoftusa.com/Smoothed-Particle-Hydrodynamics.html, accessed 15 January 2023.

[2] S. Koshizuka and Y. Oka, "Moving-Particle Semi-Implicit Method for Fragmentation of Incompressible Fluid," *Nuclear Science and Engineering*, vol. 123, no. 3, pp. 421–434, Jul. 1996, doi: doi.org/10.13182/nse96-a24205.

[3] *Particleworks 8.0 Theory Manual*, Prometech Software Inc., 2023