Adaptive particle refinement for compressible smoothed particle hydrodynamics

Daniel J. Price
School of Physics and Astronomy
Monash University
Clayton, Vic 3800, Australia
daniel.price@monash.edu

Rebecca Nealon
Department of Physics
University of Warwick
Gibbet Hill Rd, CV4 7AL, Coventry, UK
rebecca.nealon@warwick.ac.uk

I. Introduction

We present a method for splitting and merging particles in SPH based on fixed spatial volumes that is suitable for use in highly compressible flow in the public SPH code PHANTOM [1]. We refer to this method as Adaptive Particle Refinement, being conceptually similar to Adaptive Mesh Refinement (AMR) utilised in Eulerian codes [2]. The topic is not a new one for readers of SPHERIC proceedings, being a mature technique used in simulations of incompressible flow [3]–[7]. Much previous work also already exists on splitting and merging particles also for compressible flow, starting with Monaghan & Varnas in 1988 [8] and developed by many others [9]–[13]. Regular SPHERIC attendees may also be familiar with the Regularised-SPH (RSPH) technique developed over a number of years by Steinar Børve and colleagues [14]–[17] where the smoothing length was set to be piecewise-constant in steps of 2 with contributions from neighbours interpolated from a grid.

We decided to revisit this topic based on our desire for a practical implementation of APR for the public and open source astrophysical SPH code PHANTOM [1]. We thus tried to take the best of existing methods and adapt them for use in the applications for which PHANTOM is typically applied. In particular we focussed on two application areas i) circumbinary disc simulations, where one desires to resolve low density flow inside the disc cavity (building on recent efforts by [13]) and ii) planet-disc interaction simulations where one desires to better resolve the flow of material in the immediate vicinity of the planet.

A key limitation we found to existing approaches, is that although they discuss merging ‘in principle’, they did not allow merging in practice, finding it too computationally expensive [4]. This kind of one-way splitting-only refinement is relatively straightforward, and the procedure for splitting particles into children — and relaxing them into place if desired — relatively well established [3]–[7], [10]–[12]. But we desired a two-way scheme so that particles could flow through a refinement region and be successfully derefined on exit. We also follow the RSPH idea of not allowing particles of different masses to mix within the same volume, since there are known problems with mixing unequal mass particles in SPH (e.g. [6]).

II. Methods

Figure 1 illustrates the basic procedure. Based on our desire to merge and to minimise noise introduced by splitting, we refine by splitting particles into two children once the refinement boundary is crossed. To simplify matters we adopted fixed spherical volumes for our refinement regions, which can either be fixed in space or co-moving with either an SPH particle or point mass. We place the two child particles in a line perpendicular to the normal vector to the refinement boundary once the parent has crossed the boundary. Adopting a spherical boundary means finding the normal vector is simple.

For merging, we first construct a binary kd-tree where the particle distribution is progressively divided along the centre of mass, cutting along the longest spatial direction until a maximum of two particles occupies each leaf node. If the centre of mass of a leaf node crosses the derefinement boundary then the two children are merged into a parent particle that is placed at said centre of mass, with velocity and other quantities taken as a simple average from the two children.

An arbitrary number of nested refinement zones are possible, but we restrict the mass difference between adjacent levels to a factor of two.

In the course of our implementation we experimented with varying the number of children, varying the separation of child particles when first split, and a procedure for systematically relaxing the particle distribution after a split (see [18] for details). We found best results in practice with two children. We experimented with numerous relaxation procedures but in the end found it to be too expensive for practical use.

The overall scheme is described in detail in [18] but can be summarised with the following rules (from [18]): The main points of our implementation can be captured with five ‘rules’:

1) We assign all particles a refinement level, ℓ, determined from the spatial coordinates of the particle. The refinement level represents the number of refinements above the base resolution.
2) When a particle enters a new refinement zone with a given refinement level it is split if the particle’s current refinement level is less than that of the zone, or merged if the particle’s current refinement level is greater than that of the zone.

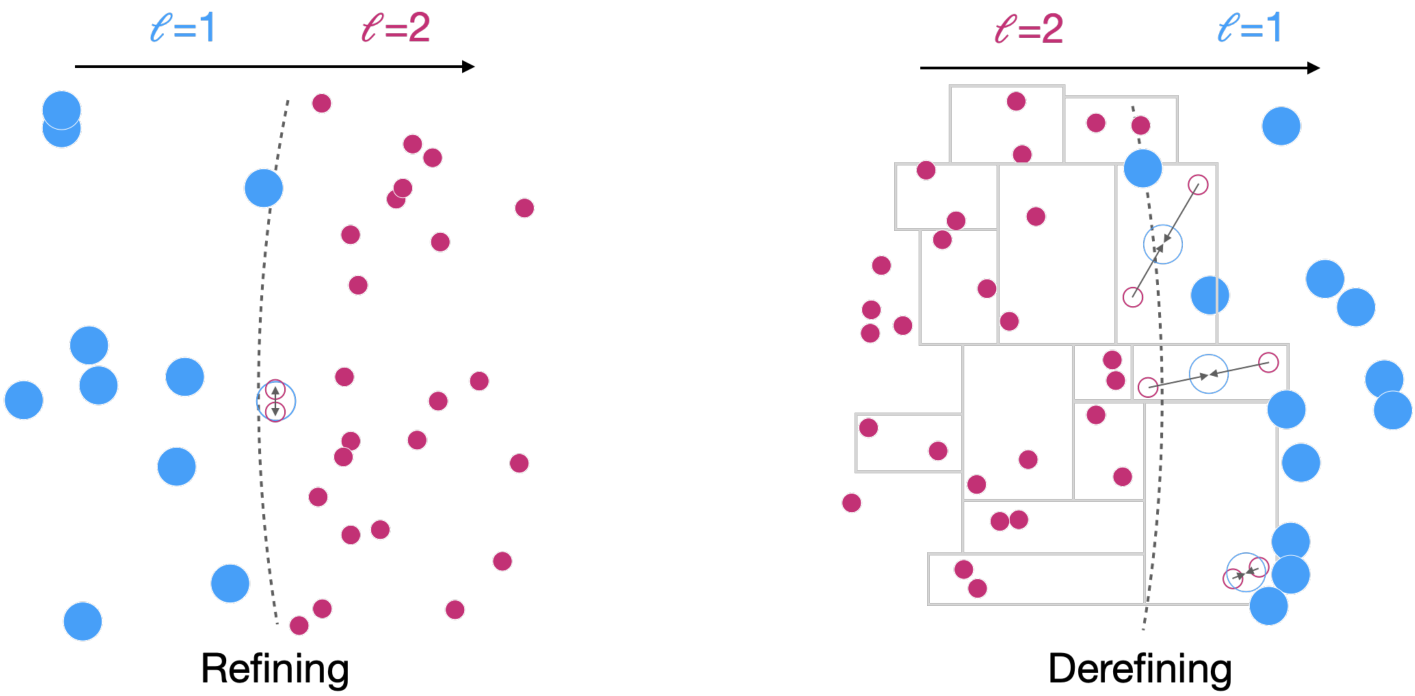

Fig. 1. Schematic of particle splitting/merging procedure. Each parent particle (blue) is split into two children (red) on crossing the refinement boundary, and the two nearest particles are merged when their centre of mass crosses the derefinement boundary

3) In simulations aiming to locally increase the resolution the default level is set to $\ell = 0$. In simulations aiming to locally decrease the resolution the default refinement level is instead set to $\ell = \ell_{\rm max}$.
4) We restrict the difference in mass between adjacent refinement levels to be strictly a factor of two.
5) As PHANTOM stores the mass by particle type, the refinement level is also used to relate the refined mass $m_p(\ell)$ from the largest particle mass $m_p(0)$ with the following relation:

$$m_p(\ell) = \frac{m_p(0)}{2^{\ell-1}}, \quad (1)$$

Although the boundary for where merging occurs is well defined, because we merge based on the child particle distribution in practice the boundary is fuzzy, with size typically of order 10% of the larger smoothing length.

III. Results and Discussion

Figure 2 shows the scheme in action. The left panel shows a low resolution simulation of a circumbinary accretion disc (where material orbits and accretes onto a binary star system) taken from [19] using 1 million particles, while the right panel shows a high resolution simulation using 8 million particles. The middle panel shows the result with three levels of APR, hence locally matching the high resolution simulation at the highest level of refinement. We also found that this simulation quantitatively matches the high resolution simulation in terms of the total mass accreted onto the two stars, with a fractional error of 1×10^{-3} in the measured mass accretion rate between the APR simulation and the high resolution simulation.

While this particular example only produces speedups of up to a factor of 2 for the APR simulation compared to running the high resolution simulation, this is because the high resolution region also corresponds to the region with the shortest timesteps (according to the local Courant condition $\Delta t \lesssim 0.3h/c_{\rm s}$). In other cases, for example in a simulation of a stellar flyby where one desires to better resolve material captured by the perturbing star, we found substantial speedups of around a factor of 6 using APR compared to the cost of performing a globally high resolution simulation.

Kinetic energy and both linear and angular momentum are conserved with APR when the children particles inherit the velocity of the parent [3], [20]. We found one disadvantage of using a relaxation procedure is that children may not necessarily be placed symmetrically around the parent particle, which affects the conservation of angular momentum.

The main limitation of our method is a ∼5% blip in density that occurs at the boundary of where particles are split or merged. This blip is not obvious in 3D simulations, but becomes obvious in tests we performed on small amplitude linear sound waves (see [18]). While we tried various methods from the literature [5]–[7], [21] to mitigate this, we found none of them satisfactory for highly compressible flow, and thus decided on limiting the number of child particles to 2 as the best strategy.

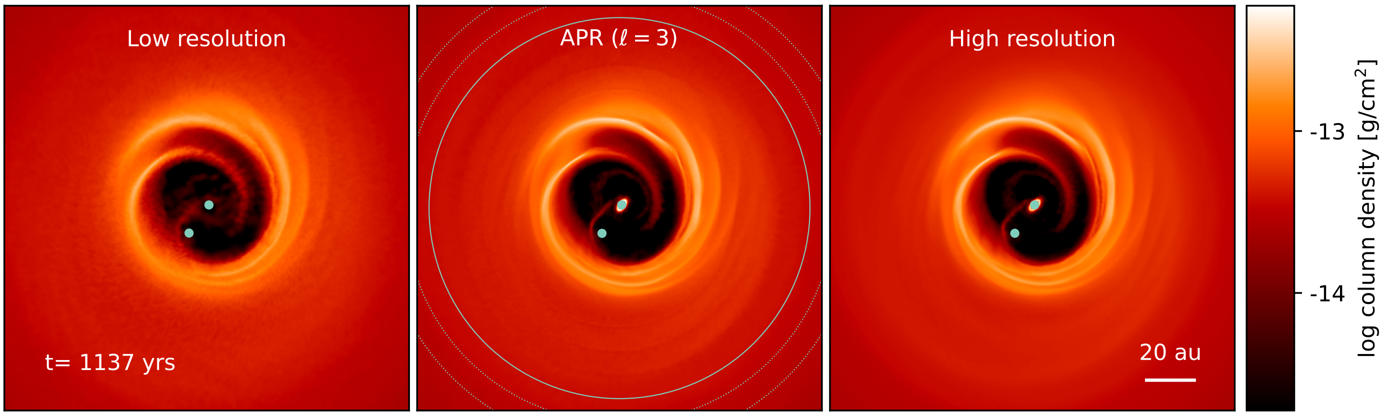

Fig. 2. Results of a circumbinary disc simulation, comparing low resolution (left) to APR (middle) to high resolution (right). The APR scheme can be seen to closely match the results of the high resolution simulation, at a fraction of the computational cost.

IV. Conclusion

In summary, we developed a new adaptive particle refinement implementation suitable for use in the PHANTOM astrophysical smoothed particle hydrodynamics code. While further details may be found in [18], we found the method to be accurate, fast and low storage, and an effective way to locally increase the resolution of simulations at low cost. It is now available in the public code.

Acknowledgment

We thank James Wurster for significant contributions in the early stages of this work, and thank Matthew Bate, Richard Booth, Terry Tricco, Alessia Franchini and Sahl Rowther for useful discussions. R.N. acknowledges support from UKRI/EPSRC through a Stephen Hawking Fellowship (EP/T017287/1). This work was performed using Avon, the HPC clusters at the University of Warwick. DP is grateful for Australian Research Council funding via DP220103767 and DP240103290. We thank the Monash-Warwick alliance for sponsorship of the PHANTOM users workshop where this work commenced.

References

[1] D. Price, J. Wurster, T. S. Tricco, C. Nixon, S. Toupin, A. Pettitt *et al.*, "Phantom: A Smoothed Particle Hydrodynamics and Magnetohydrodynamics Code for Astrophysics," PASA, vol. 35, p. e031, Sep. 2018.

[2] M. J. Berger and P. Colella, "Local Adaptive Mesh Refinement for Shock Hydrodynamics," *J. Comp. Phys.*, vol. 82, no. 1, pp. 64–84, May 1989.

[3] J. Feldman and J. Bonet, "Dynamic refinement and boundary contact forces in SPH with applications in fluid flow problems," *International Journal for Numerical Methods in Engineering*, vol. 72, no. 3, pp. 295–324, Oct. 2007.

[4] R. Vacondio, B. D. Rogers, P. K. Stansby, P. Mignosa, and J. Feldman, "Variable resolution for SPH: A dynamic particle coalescing and splitting scheme," *Computer Methods in Applied Mechanics and Engineering*, vol. 256, pp. 132–148, Apr. 2013.

[5] D. Barcarolo, D. Le Touzé, G. Oger, and F. de Vuyst, "Adaptive particle refinement and derefinement applied to the smoothed particle hydrodynamics method," *J. Comp. Phys.*, vol. 273, pp. 640–657, 2014. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0021999114004033

[6] L. Chiron, G. Oger, M. de Leffe, and D. Le Touzé, "Analysis and improvements of adaptive particle refinement (apr) through cpu time, accuracy and robustness considerations," *J. Comp. Phys.*, vol. 354, pp. 552–575, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0021999117308082

[7] T. Gao, H. Qiu, and L. Fu, "A block-based adaptive particle refinement sph method for fluid–structure interaction problems," *Computer Methods in Applied Mechanics and Engineering*, vol. 399, p. 115356, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0045782522004376

[8] J. J. Monaghan and S. R. Varnas, "The dynamics of interstellar cloud complexes," MNRAS, vol. 231, pp. 515–534, Mar. 1988.

[9] Z. Meglicki, D. Wickramasinghe, and G. V. Bicknell, "Three-Dimensional Structure of Truncated Accretion Discs in Close Binaries," MNRAS, vol. 264, p. 691, Oct. 1993.

[10] S. Kitsionas and A. P. Whitworth, "Smoothed Particle Hydrodynamics with particle splitting, applied to self-gravitating collapse," MNRAS, vol. 330, no. 1, pp. 129–136, Feb. 2002.

[11] ——, "High-resolution simulations of clump-clump collisions using SPH with particle splitting," MNRAS, vol. 378, no. 2, pp. 507–524, Jun. 2007.

[12] D. Anglés-Alcázar *et al.*, "Cosmological Simulations of Quasar Fueling to Subparsec Scales Using Lagrangian Hyper-refinement," ApJ, vol. 917, no. 2, p. 53, Aug. 2021.

[13] A. Franchini, A. Lupi, and A. Sesana, "Resolving Massive Black Hole Binary Evolution via Adaptive Particle Splitting," *ApJL*, vol. 929, no. 1, p. L13, Apr. 2022.

[14] S. Børve, M. Omang, and J. Trulsen, "Regularized Smoothed Particle Hydrodynamics: A New Approach to Simulating Magnetohydrodynamic Shocks," ApJ, vol. 561, no. 1, pp. 82–93, Nov. 2001.

[15] ——, "Regularized smoothed particle hydrodynamics with improved multi-resolution handling," *J. Comp. Phys.*, vol. 208, no. 1, pp. 345–367, Sep. 2005.

[16] ——, "Multidimensional MHD Shock Tests of Regularized Smoothed Particle Hydrodynamics," ApJ, vol. 652, no. 2, pp. 1306–1317, Dec. 2006.

[17] S. Børve, R. Speith, and J. Trulsen, "Numerical Dissipation in RSPH Simulations of Astrophysical Flows with Application to Protoplanetary Disks," ApJ, vol. 701, no. 2, pp. 1269–1282, Aug. 2009.

[18] R. Nealon and D. Price, "Adaptive particle refinement for compressible SPH," *accepted to PASA*, p. arXiv:2409.11470, 2025.

[19] D. Price *et al.*, "Circumbinary, not transitional: on the spiral arms, cavity, shadows, fast radial flows, streamers, and horseshoe in the HD 142527 disc," MNRAS, vol. 477, pp. 1270–1284, Jun. 2018.

[20] Y. López, D. Roose, and C. Recarey Morfa, "Dynamic particle refinement in sph: Application to free surface flow and non-cohesive soil simulations," *Computational Mechanics*, vol. 51, 05 2013.

[21] S. Lind, R. Xu, P. Stansby, and B. Rogers, "Incompressible sph for free-surface flows: A generalised diffusion-based algorithm for stability and validations for impulsive flows and propagating waves," *J. Comp. Phys.*, vol. 231, no. 4, pp. 1499–1523, 2012. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0021999111006279