Planetary Giant Impacts: Convergence of High-Resolution Simulations using Efficient Spherical Initial Conditions and SWIFT

J. A. Kegerreis¹1, V. R. Eke¹, P. Gonnet², D. G. Korycansky³, R. J. Massey¹, M. Schaller⁴, L. F. A. Teodoro⁵ ¹Institute for Computational Cosmology, Durham University, Durham, DH1 3LE, UK

²Google AI Perception, Google Switzerland, 8002 Zurich, Switzerland

³CODEP, Earth Sciences, University of California, Santa Cruz, CA, USA

⁴Leiden Observatory, Niels Bohrweg 2, 2333 CA Leiden, Netherlands

⁵BAERI/NASA Ames Research Center, Moffett Field, CA, USA

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

We perform simulations of giant impacts onto the young Uranus using smoothed particle hydrodynamics (SPH) with over 100 million particles. This $100-1000 \times$ improvement in particle number reveals that simulations with below 10^7 particles fail to converge on even bulk properties like the post-impact rotation period, or on the detailed erosion of the atmosphere. Higher resolutions appear to determine these large-scale results reliably, but even 10^8 particles may not be sufficient to study the detailed composition of the debris - finding that almost an order of magnitude more rock is ejected beyond the Roche radius than with 10⁵ particles. We present two software developments that enable this increase in the feasible number of particles. First, we present an algorithm to place any number of particles in a spherical shell such that they all have an SPH density within 1% of the desired value. Particles in model planets built from these nested shells have a root-mean-squared velocity below 1% of the escape speed, which avoids the need for long precursor simulations to produce relaxed initial conditions. Second, we develop the hydrodynamics code SWIFT for planetary simulations. SWIFT uses task-based parallelism and other modern algorithmic approaches to take full advantage of contemporary supercomputer architectures. Both the particle placement code and SWIFT are publicly released.

> Key words: methods: numerical - hydrodynamics - planets and satellites: physical evolution – planets and satellites: atmosphere

INTRODUCTION 1

Giant impacts are thought to dominate many planets' late accretion and evolution. We see the consequences of these violent events on almost every planet in our solar system, from the formation of Earth's Moon to the odd obliquity of Uranus spinning on its side. As such, they are expected to play a similarly important role in the evolution of the many exoplanetary systems that are now being observed in detail. These complicated and highly non-linear processes are most commonly studied using smoothed particle hydrodynamics (SPH) simulations (e.g. Benz et al. 1986).

SPH is a Lagrangian (particle-based) method used in a wide range of topics in astrophysics and many other fields, from planetary impacts and supernovae to galaxy evolution and cosmology (Springel 2010; Monaghan 2012). As well as correctly evolving the simulation particles with time, it is crucial to start from appropriate initial conditions for any model's evolution to accurately reflect its real-world counterpart. Furthermore, enough particles must be used to resolve the physical processes in sufficient detail, and recent work has shown that standard-resolution simulations (10^5) to 10^6 particles) can produce unreliable results that have not converged numerically (Hosono et al. 2017; Genda et al. 2015). This motivates the pursuit of simulation codes that can take full advantage of contemporary supercomputing architectures, enabling more particles to be used to run suitable convergence tests and, hopefully, simulations with sufficiently high resolution.

2 J. A. Kegerreis et al.

Towards this end, we present the simple SEA¹ scheme for creating optimal spherical arrangements of particles (§2.1, §3.1) and the hydrodynamics code SWIFT² that we have developed to run planetary impact (and cosmological) simulations (§2.2). We use these tools to model giant impacts onto a young Uranus at high resolution using over 10^8 SPH particles, and test the convergence of various physical properties with increasing particle number (§1.2, §3.2). We then present conclusions in §4.

1.1 Particle Placement and Initial Conditions

Many problems in astrophysics feature spherical symmetry, such as those involving stars or planets. Before one can simulate and study these problems with a particle-based method like SPH, each initial object must first be converted into an appropriate set of particles. Two common approaches to creating arrangements of SPH particles in spheres are: (1) to use a lattice that can be distorted until it approximately matches the required shape and densities; and (2) to relax an imperfect initial state into a fully settled one with a preproduction simulation.

A third, more recent approach is to arrange the particles analytically while accounting for the spherical symmetry from the outset, by placing particles in nested spherical shells (Saff & Kuijlaars 1997; Raskin & Owen 2016; Reinhardt & Stadel 2017). These methods aim to combine the minimal computation required for lattice methods with the settled and symmetric properties of simulated glasses. We present a comparable scheme that further ensures every particle's SPH density is within 1% of the desired value. This leads to initial conditions that are quick and simple to produce, close to equilibrium, and in which every particle has a realistic density and, therefore, pressure.

Lattice-based methods are popular because they are easy to implement and, since the inter-particle separations are uniform by construction, they can accurately match a simple density profile. This can be achieved either by stretching the lattice radially or by varying the particle mass – although keeping the masses of all particles very similar is usually desirable. However, the grid-like properties of a lattice introduce unwanted anisotropies to a problem and may be unstable to perturbations (Herant 1994; Morris 1996; Lombardi et al. 1999).

Furthermore, a spherically symmetric object like a planet or star features important boundaries at specific radii. Both the outer surface and internal layers require discontinuities in density and material. The particles in a lattice are dispersed at all radii, so cannot reproduce such sharp changes at these boundaries. A similarly quick and simple alternative to lattice methods is to randomly place particles following an appropriate probability distribution function, either restricted in nested shells or anywhere in the sphere. However, these methods are noisy and result in extreme variations in local particle densities. In SPH, the density of a particle is estimated by summing the masses of typically ~50 nearby 'neighbour' particles, weighting by a 3D-Gaussian-shaped kernel that decreases the contribution of more-distant neighbours. Thus a particle that is placed too close to another will have a higher density and not be in equilibrium. The accuracy of every particle's density is important because of how stiff the equation of state (EoS) can be for a material, such as the granite planetary example we test here. This means that a slightly too-dense particle will be assigned a dramatically too-high pressure by the EoS, leading to unphysical behaviour as soon as the simulation is started. In the case of a tabulated EoS, this may also cause practical problems by pushing a particle outside of the parameter space covered by the tables.

An obvious improvement on these crude analytical distribution methods is to run a simulation that iterates the initial particle positions towards a more stable state. One approach is to use an inverse gravitational field to repel the particles from each other (Wang & White 2007). A more sophisticated version of this was developed by Diehl et al. (2015) based on weighted Voronoi tessellations. Another method is to add a damping force to reduce any transient velocities as the particles are allowed to evolve under otherwise-normal gravitational and material pressure forces. In all cases, the simulation is run until a condition is met to call the system 'relaxed', such as when the particle velocities or accelerations reach some small value.

These methods can generate particle configurations that are stable and relaxed, but at a cost of performing extra simulations. Especially for large numbers of particles, this can be a computationally expensive process and can take large amounts of time, comparable to the final simulation for which the initial conditions are being generated. Depending on the method used, the particles may also settle to a distribution somewhat different to the desired initial profile.

The spherical symmetry and sharp radial boundaries of astrophysical objects strongly motivate the arrangement of particles in nested spherical shells. If the particles could be distributed uniformly in each shell, then no computationally expensive simulation would be required to create relaxed initial conditions. However, the equidistant distribution of points on the surface of a sphere is a challenging problem, and has been studied for applications in a wide variety of fields: from finding stable molecular structures like buckminsterfullerene to making area-integral approximations, in addition to the pure mathematical curiosity of such a trivial question in 2D (equally spaced points on a circle) becoming so complicated in higher dimensions (Saff & Kuijlaars 1997).

Similar ideas motivated the work of Raskin & Owen (2016) and Reinhardt & Stadel (2017), who both presented algorithms for arranging particles in spherical shells. One issue with Raskin & Owen (2016)'s method is that in each shell there are a few particles with large overdensities, placing the particles slightly out of equilibrium (see §3.1). Reinhardt & Stadel (2017) divide the sphere into equal regions that can be further subdivided (using the HEALPix scheme), with the disadvantage that only sparsely distributed numbers of particles $(12 \times 4^n \text{ for } n \in \mathbb{N})$ can be placed in each shell. Furthermore, some particles in each shell show

¹ The SEAGen code is publicly available at github.com/jkeger/seagen and the python module seagen can be installed directly with pip.

 $^{^2}$ SWIFT is in open development and publicly available at swift.dur.ac.uk.

SPH densities more than 5% discrepant from the desired profile density (their Fig. 4).

In §2.1, we present an algorithm for arranging any number of particles in a spherical shell such that every particle has an SPH density within 1% of the median. Our method involves a simple division of the sphere into equal-area regions arranged in latitudinal collars, followed by slightly stretching the collars away from the poles. Concentric shells can then be set up to precisely follow an arbitrary radial density profile, taking care to align the shells with any radial boundaries. We apply this stretched equal-area (SEA) algorithm to create near-equilibrium models of planets, and present the results in §3.1.

1.2 Convergence and Uranus Giant Impacts

The need to increase resolution to improve studies of existing topics was recently demonstrated by Hosono et al. (2017). Concerningly, they found giant impact simulations that gave apparently reliable results with up to 10^6 particles had not actually converged when re-tested with 10^7-10^8 . Genda et al. (2015) also found incomplete convergence of disruptive impact simulations with up to 5×10^6 particles.

A numerically converged result is not necessarily physically correct. For example, several studies (e.g. Woolfson 2007; Deng et al. 2017) have pointed out the difficulties for SPH in modelling the interaction of multiple materials or the treatment of density discontinuities, which may not be immediately fixed by higher resolutions. That said, it is crucial that we at least obtain a reliable answer to the (imperfect or not) question that we ask the computer to solve, so convergence is an important first step.

As an example with which to investigate convergence and test the simulation tools presented in this paper, we consider the giant impact that likely knocked over the planet Uranus to spin on its side. Previously, we ran SPH simulations to study the consequences of this violent event using $\sim 10^6$ particles (Kegerreis et al. 2018, hereafter K18) – as an improvement on the $<10^4$ particles in the single previous study by Slattery et al. (1992) almost 30 years ago. As well as confirming that the impact can explain Uranus' spin, we found that with a grazing collision the impactor could form a thin shell around the planet's ice layer, perhaps trapping the interior heat to help explain the freezing outer temperatures. $\sim 10\%$ of the target's atmosphere becomes unbound to escape from the system and a small amount of the impactor's rocky core is ejected into the debris disk. Kurosaki & Inutsuka (2019) recently explored a different, complementary part of the wide parameter space with $\sim 10^5$ SPH particle simulations. They varied the entropy of the proto-Uranus target to examine the effects on the angular momentum and the debris.

In §2.2 we summarise the SWIFT hydrodynamics code and its development to run these planetary simulations and take advantage of contemporary supercomputer architectures. In §3.2 we use SWIFT and the SEA particle placement method to repeat simulations of Uranus giant impacts from K18 using 10^5 up to 10^8 SPH particles, and test the convergence of the post-impact planet's rotation rate, the erosion of the atmosphere, and the ejection of rocky material into the debris disk.

2 METHODS

2.1 Particle Placement and Initial Conditions

The goal is to distribute a number of similar-mass particles in a sphere, such that the SPH density of every particle accurately matches a given density profile (see §A). In order to follow an arbitrary radial profile that may include sharp discontinuities, such as a core-mantle boundary or a planet's surface, it is convenient to distribute the particles in spherical shells. The particles can then be assigned any property using other radial profiles, such as their material type and temperature or internal energy.

The two inputs for this problem are the desired total number of particles and the radial density profile. The profile is first used to find the enclosed mass at each radius. The number of particles then gives the nominal particle mass. We iterate outwards from the centre, placing particles in successive shells, following the density profile. First, we must determine the radius of each shell and how many particles are required to account for its mass (§2.1.1). Then, the question is how to arrange an arbitrary number of particles on a spherical shell, for which we describe our stretched equalarea (SEA) method (§2.1.2).

2.1.1 Shells and Layers

We begin by placing a tetrahedron of particles near the centre, so the first 'shell' is actually the sphere that encloses the mass of four particles. If this central sphere has radius dr_c and density ρ_c , then the thickness, dr, of a subsequent shell with density ρ is

$$dr = dr_c \left(\frac{\rho_c}{\rho}\right)^{1/3} \,. \tag{1}$$

The number of particles in a shell is then simply the mass of that shell divided by the nominal particle mass. This must be rounded to an integer, giving an actual particle mass in each shell that may be slightly different to the nominal mass. This amounts to maximum deviations of ~1% for 10^6 total particles and ~0.1% for 10^8 . The shell thickness could be tweaked instead to enforce strictly equal particle masses. The particles in the shell are then all assigned the same properties (e.g. temperature), set by the mass-weighted mean of the profile values across the shell.

It is important to note that this shell spacing will, in general, lead to shell boundaries that do not line up with any boundaries in the profile – whether simply the outer profile edge or internal boundaries separating layers inside a planet or star. In the first case of a single-layer profile, the penultimate particle shell will typically end close to the outer edge. This leaves a thin and low-mass outermost shell with only a small number of particles that both cannot adequately cover the large area and will be too close in radius to the previous shell. For interior boundaries such as between core and mantle layers, a shell will typically straddle the discontinuity. The particles in this shell then try in vain to represent some of both materials and conditions.

To avoid these problems, we slightly tweak the input particle mass to change the mass of the first core shell and hence its radius. This influences the radii of all the shells (Eqn. 1). We iterate the input particle masses until the boundary of the outermost shell in the first (or only) layer

4 J. A. Kegerreis et al.

coincides with the profile's boundary. This leads to a slightly different total number of particles as well, but ensures a proper particle representation of the final shell in this layer, as well as of the first shell of the next layer if there is one.

A similar issue and solution arises for any subsequent boundaries. To maintain a similar particle mass in all layers, we do not change the particle mass again. Instead, we tweak the number of particles in the first shell of each outer layer. This changes the mass of that shell and hence its radius, as before. By using the thickness and density of this shell in Eqn. 1 instead of the central shell, this leads to appropriate changes for all the shells in this layer. We iterate over slightly different numbers of particles in the first shell until the outermost shell's boundary coincides with the profile boundary of this layer. This is repeated at the start of each layer until a particle shell boundary matches every profile boundary both internally and at the profile's edge.

One remaining decision is at what radius to place the particles within each shell. Two average radii to consider are $r_{1/2}$, half-way between the inner and outer radii of the shell, and $r_{\rm m-w}$, the mass-weighted mean radius. For a slowly changing density profile and/or many particles that lead to thin shells, the density is roughly constant throughout the shell and $r_{\rm m-w} > r_{1/2}$ because the mass increases with $4\pi r^2$. In the vast majority of shells, where $dr \ll r$, these two radii are approximately equal. However, at small radii near the core, placing the particles at $r_{1/2}$ results in too-high densities, and $r_{\rm m-w}$ gives too-low densities. We found that placing the particles at $\frac{1}{2}(r_{1/2} + r_{\rm m-w})$ correctly matches the mean SPH density of the particles in each shell to the profile density at that radius.

2.1.2 Particles on a Sphere

For every shell, we now have a number of particles, N, to distribute on the surface of a sphere. We begin by considering the division of a (unit) sphere into equal-area regions with small diameters, following the algorithm described by Leopardi (2007) with minor modifications. The particles can then be placed in the centre of each region.

We further impose a stretching of the regions by latitude, to improve the particle density near the poles. Finally, each shell is randomly rotated so that the particles at the poles do not line up in successive shells.

For comparison, we also test the recursive primitive refinement and parametrised spiral (RPR+PS) method described by Raskin & Owen (2016). Their method uses subdivisions of the Platonic solids for low-N shells and a spiral placement algorithm for larger numbers of particles.

2.1.3 Equal-Area Regions

This method is also illustrated in Fig. 1 and a finished example with N=100 is shown in Fig. 2.

For N regions on a sphere, the area of each one will be

$$A_{\rm reg} = 4\pi/N \ . \tag{2}$$

The bounding colatitude of a polar cap with area A_{cap} is

$$\theta = 2 \arcsin\left(\sqrt{\frac{A_{\rm cap}}{4\pi}}\right) \,, \tag{3}$$



Figure 1. An example division of a sphere into 20 equal regions, demonstrating the main steps in the algorithm: (a) set the polar caps and the initial collar latitudes; (b) tweak the collar latitudes so that they each contain an integer number of regions; (c) divide each collar into equal regions; (d) rotate the collars to maximise the minimum separation of adjacent regions.

which for $A_{\rm cap} = A_{\rm reg}$ gives the colatitude of the singleregion north pole cap, $\theta_{\rm cap}$, and south pole cap, $\pi - \theta_{\rm cap}$.

We start by dividing the rest of the sphere (between the two polar caps) into collars with ideal initial heights of $\sqrt{A_{\text{reg}}}$. This gives the number of collars (when rounded to an integer),

$$N_{\rm col} = {\rm round} \left[\frac{\pi - 2\theta_{\rm cap}}{\sqrt{A_{\rm reg}}} \right] \,, \tag{4}$$

and the actual initial collar height,

$$\theta_{\rm col} = \left(\frac{\pi - 2\theta_{\rm cap}}{N_{\rm col}}\right) \,, \tag{5}$$

(Fig. 1a). We then divide each initial collar i into the closest integer number of regions. The area of each collar is

$$A_{i} = 4\pi \left(\sin^{2} \left(\frac{\theta_{i}}{2} \right) - \sin^{2} \left(\frac{\theta_{i-1}}{2} \right) \right) , \qquad (6)$$

so the ideal number of regions in each collar i is

$$N_i' = \frac{A_i}{A_{\rm reg}} \ . \tag{7}$$

This must be rounded to the actual integer number of regions, N_i . The cumulative discrepancy, d_i , from the ideal number of regions must be included to ensure that the total number of regions is unchanged:

$$N_i = \operatorname{round} \left[N'_i + d_i \right] \tag{8}$$

$$d_{i+1} = d_i + N'_i - N_i . (9)$$

Starting from the north pole and using the cumulative number of regions in each collar, $N_{\leq i}$, we find the final colatitude of each collar by calculating the colatitude of the cap that contains the same area as $N_{\leq i}$ regions:

$$\theta_i = 2 \arcsin\left(\sqrt{\frac{N_{\le i} A_{\text{reg}}}{4\pi}}\right) \,, \tag{10}$$



Figure 2. An example of 100 particles distributed on a sphere using the SEA (equal-area and subsequent latitude-stretching) method. The colours highlight each collar of particles. The SPH densities of these particles are shown by the purple points in Fig. 3.

where i = 1 is the north pole cap (Fig. 1b).

The points in the centre of each region j in collar i then have

$$\theta = \frac{1}{2} \left(\theta_i + \theta_{i+1} \right) \tag{11}$$

$$\phi = \phi_0 + j \,\Delta\phi_i \;, \tag{12}$$

where ϕ_0 is the starting longitude and $\Delta \phi_i = \frac{2\pi}{N_{\leq i}}$ is the angle between adjacent points (Fig. 1c).

We choose the starting longitude of each collar, ϕ_0 , to maximise the minimum separation between the points on adjacent collars (Fig. 1d). This helps to prevent local overdensities. If N_i and N_{i-1} are both odd or both even, then ϕ_0 is half the smaller of $\Delta \phi_i$ and $\Delta \phi_{i-1}$. If one is odd and the other is even, then ϕ_0 must be half of the even one's $\Delta \phi$, to prevent two particles in adjacent collars from having the same ϕ and being too close together.

Finally, ϕ_0 should be additionally offset by $m \Delta \phi_{i-1}$, where *m* is a random integer between 0 and N_{i-1} . Thus, the ϕ_0 rotation will be with respect to a random particle in the previous collar. This prevents successive collars with large N_i (and hence small ϕ_0) from creating a sequence of nearly adjacent particles in successive collars.

2.1.4 Latitude Stretching

The equal-area scheme described in §2.1.3 results in a small local overdensity of particles near the poles. We can make the particle density more uniform by stretching the collars near the poles. However, the collars near the equator must not be overly squashed. Therefore, the (absolute) latitude of each point, $|\frac{\pi}{2} - \theta|$, should be reduced by an amount that varies with latitude, from maximum stretching at the poles to 0 at the equator. Of course, the size of the shift at all latitudes depends on the initial size of the collars, which is set by the total number of particles. The collar height and the required shift will decrease in proportion with the square root of the number of particles. Thus, the appropriate stretching can be

given by:

$$\theta' = \theta + \left(\frac{\pi}{2} - \theta\right) \times aN^{-\frac{1}{2}} \exp\left[-\frac{\frac{\pi}{2} - |\frac{\pi}{2} - \theta|}{\pi bN^{-\frac{1}{2}}}\right], \quad (13)$$

where a = 0.2 and b = 2 (tested for $80 \le N \le 10^6$). For N < 80, we fit *a* and *b* manually to ensure that the maximum deviation of any particle's density from the mean is less than $\pm 1\%$. This requires *a* to vary (non-monotonically) between 0.18 and 0.27, with *b* following this variation as b = 10 a, and is only relevant for the innermost one or two lowest mass shells.

2.2 Planetary Simulations with SWIFT

SWIFT (SPH With Inter-dependent Fine-grained Tasking) is a hydrodynamics and gravity code for astrophysics and cosmology in open development (swift.dur.ac.uk), designed from the ground up to run fast and scale well on shared/distributed-memory architectures (Schaller et al. 2016).

For the past decade, physical limitations have kept the speed of individual processor cores constrained, so instead of getting faster, supercomputers are getting more parallel. This makes it ever more important to share the work evenly between every part of the computer so that no processors are sitting idle and wasting time.

SWIFT can function as a drop-in replacement for the Gadget-2 code, which has been widely used for cosmological and planetary impact simulations (Springel 2005; \acute{C} uk & Stewart 2012), but with a >30× faster runtime on representative cosmological problems (Borrow et al. 2018). This speed is partly a result of SWIFT's task-based approach to parallelism and domain decomposition for the gravity and SPH calculations (Gonnet 2015). By evaluating and dividing up the work instead of just the data, this provides a dynamic way to achieve good load balancing across multiple processors within a shared-memory node. The tasks are decomposed over the network in distributed memory systems using a graph-partitioning algorithm, weighting each task by the estimated computational work it requires. Combined with using asynchronous communications that are themselves treated as normal tasks, this allows the code to scale well (Schaller et al. 2016). Core routines, including the direct interaction between particles, have then been optimized using vector instructions (Willis et al. 2018).

In some respects, giant impact simulations pose a harder challenge for load balancing than the cosmological simulations that SWIFT is also designed for. For a large patch of the universe, although the density becomes very much higher in a galaxy than a void, the local average density is roughly constant across a simulation box. Even a crude division of particles by position in the box to different computing cores can somewhat effectively speed up the calculation, and a more careful decomposition like SWIFT's can produce excellent strong scaling across hundreds of thousands of cores (Borrow et al. 2018).

In contrast, for a giant impact, almost all the mass (and hence particles) is in the planet at the centre. If we use a large simulation box in order to follow the ejected debris, then the vast majority of particles can easily occupy less than 0.01% of the volume. This is similar to cosmological

'zoom-in' simulations that use a high-resolution region to focus on a single galaxy or halo. This firstly makes it harder to divide up particles between computing nodes, and secondly can require much more frequent communication. This makes it much less efficient to use a large numbers of cores, and difficult to fully utilise a large supercomputer to run a single planetary simulation very quickly.

Happily, most studies of giant impacts can be reframed as 'embarrassingly parallel' problems because, instead of investigating one specific collision in extreme detail, the usual aim is to study a wide range of scenarios, such as varying the impact angle and speed. For this reason, perfect scaling across many distributed-memory nodes or MPI ranks is not as important. Many impacts can each be simulated on their own single (or small number of) shared-memory node(s). SWIFT then uses threads and SIMD vectorisation to parallelise efficiently across the tens of cores within each node. However, as we investigate in §3.2, even for parameterspace surveys, large numbers of particles may be necessary to obtain sufficiently converged results, depending on the property being studied.

2.2.1 Planetary SPH

SWIFT has a modular structure that separates different code sections for clean modifications to, for example, the physics or the hydrodynamics scheme without affecting (or even being aware of) the parallelisation and other structural components. Any such module is switched in or out with configuration flags, allowing SWIFT to run planetary, cosmological, or any other simulation as required.

The hydrodynamics scheme used for the simulations in this paper uses a simple 'vanilla' form of SPH as described in e.g. Price (2012), with the Balsara switch for the artificial viscosity (Balsara 1995). Multiple other schemes are also implemented in SWIFT, as well as various SPH kernels. Here, we use the simple 3D cubic spline kernel with 48 neighbours, corresponding to a ratio of smoothing length to inter-particle separation of $\gamma = 1.2348$ (Dehnen & Aly 2012). The default artificial viscosity parameters for the Monaghan (1992) model are set to $\alpha = 1.5$ and $\beta = 2 \alpha$, as is typical in the literature (e.g. Reinhardt & Stadel 2017).

The equation of state (EoS) for a material relates its pressure to its density and temperature (or internal energy or entropy). So far,³ we have implemented several Tillotson, SESAME, and Hubbard & MacFarlane (1980) (for Uranus materials) EoS in SWIFT, as well as an ideal or isothermal gas. Any number of these different materials can be simulated together, as is required in a multi-layered planet, for example.

3 RESULTS AND DISCUSSION

3.1 Particle Placement

In this section, we first test the arrangement of particles on an isolated spherical shell. Then, we investigate full 3D initial conditions for a simple Earth-mass planet, consider-



Figure 3. The SPH densities of 100 particles placed using three different schemes, normalised by the median density. The grey lines show $\pm 1\%$ of the median. The 3D positions of these SEA particles are illustrated in Fig. 2.

ing the SPH densities of the particles in their initial positions and how close they are to equilibrium when allowed to evolve.

Fig. 3 shows the densities of 100 particles arranged on a unit spherical shell using three different methods: Raskin & Owen (2016)'s recursive primitive refinement and parametrised spiral method (RPR+PS, specifically PS in this case) and our equal-area method without (EA) and with (SEA) the extra latitude stretching, as described in §2.1.2.

The RPR+PS and EA methods both show significant overdensities at the poles, with maximum deviations from the median density approaching 20% and 10% respectively. This is still a big improvement on a random distribution of particles on a shell, which leads to densities that are wrong by a factor of >10. The SEA stretching reduces the scatter to less than 1%, with typical maximum deviations of 0.5%, depending on the exact number of particles. Only 100 particles are shown here for clarity; the three methods show similar relative deviations for 10^2-10^6 particles in a single shell.

Unfortunately, this dramatic improvement of SEA over the unstretched EA method cannot be replicated for RPR+PS because the distribution of particles is not azimuthally symmetric. Stretching the RPR+PS particles at the poles reduces the overdensity for some particles but creates unavoidable underdensities for others because of their asymmetry.

To investigate how these properties of an isolated shell translate into nested shells in 3D, we now consider a full model of an Earth-mass planet with $\sim 10^5$ particles (see appendix A). The results from using the same three placement methods are shown in Fig. 4. As in the isolated-shell case, the RPR+PS particles show a large range of densities, with a systematic spread of particle densities more than 10% discrepant from the profile. The unstretched EA method shows similar density discrepancies around 4%, while the

 $^{^{3}}$ The simulations in this paper used SWIFT version 0.8.1.



Figure 4. The SPH densities of $\sim 10^5$ particles placed using the three different shell schemes as labelled in the legend. The EA and RPR+PS particles are shown offset to slightly higher radii for clarity. The black line shows the input density profile, representing a simple model of an Earth-mass planet. The SEA particles' densities stay within 1% of the profile, as in the Fig. 3 isolated shell case.

SEA stretching again ensures the scatter is within 1% of the profile density. These values are for a cubic spline kernel with 48 neighbours. Using another common example of the Wendland-C6 kernel with 200 neighbours yields the same qualitative results but reduces the density scatter in all cases by roughly $\frac{1}{2}$.

The underdensity of particles in the outermost shell is caused by the nature of the SPH density calculation, so is seen equally for all methods. The spherical kernel volume extends into the empty space above the planet's surface without finding any neighbours, artificially reducing the density.

It is noteworthy that the density deviations of the RPR+PS and EA methods were reduced when switching from the 2D to the 3D case, while the SEA deviations were approximately unchanged. This reflects the contributions of the particles in other shells to the SPH density. The high overdensities are reduced in 3D because the nearby particles in adjacent shells are also summed over, mitigating the impact of the too-close particles in the same shell. For SEA, the particles in the randomly rotated adjacent shells are just as likely to be very slightly too close or too far as the particles in the same shell, so the density discrepancies are largely unchanged. This suggests that there would be little benefit to improving the distribution of particles within each shell beyond that of SEA, e.g. by running a relaxing simulation within each shell. Even if the particles in every isolated shell were perfectly arranged, then the imperfect contributions from adjacent-shell particles would negate any improvement. So, if even smaller density deviations were desired, then it would be necessary to consider all particles at once.

The actual success of our method is determined by how

close the particles are to equilibrium when allowed to evolve in a simulation. A standard criterion for initial conditions to be considered 'relaxed' enough for use is that the root mean square velocity, $v_{\rm rms}$, is below ~1% of the escape speed, here $v_{\rm esc} = 11.2$ km s⁻¹. Thanks to their precise densities, the SEA particles immediately have $v_{\rm rms}$ below 0.01 $v_{\rm esc}$, and the maximum particle speed first peaks at under 0.04 $v_{\rm esc}$. ('Immediately' here meaning the fastest speeds the particles reach, soon after being allowed to evolve from a stationary start.) In comparison, a random distribution of particles in shells has initial $v_{\rm rms} = 0.2 v_{\rm esc}$.

Most of the SEA particles' motion is caused by the previously mentioned underdensity of the outermost shell, which causes the entire planet to gently oscillate and settle into a slightly lower density profile. Because this dominates the discrepancy from an equilibrium state, the RPR+PS particles' $v_{\rm rms}$ is almost identical to SEA in spite of their comparatively noisy densities. Their maximum speed is slightly higher at 0.07 $v_{\rm esc}$. If a modified density estimator is used to fix the outer boundary problem, then a larger difference might be expected between the two methods. Planets with layers of different materials – such as the proto-Uranus and impactor in §3.2 – face similar SPH density problems at interior boundaries as well.

We confirmed that these relaxed SEA results are unchanged for Moon- and Pluto-mass planets (~0.01 and 0.002 M_{\oplus}), which are less strongly gravitationally bound, making them slightly less stable. However, the Tillotson EoS used here (Tillotson 1962) is even steeper close to the low density at which the pressure is zero, as is the case for other EoS and depending on the temperature. This exacerbates any density errors into even greater pressure discrepancies. For RPR+PS, some under-dense particles in the Plutomass planet are even pushed below the zero-pressure density, while the most over-dense ones get assigned a pressure over 4 times the desired value. Nevertheless, these particles can quickly be relaxed without much affecting the overall structure or $v_{\rm rms}$. SEA has the mild advantage that it avoids such issues in the first place, and requires similarly trivial computation to generate the initial conditions.

The SEAGen code for quickly generating both isolated shells and full spheres of points is publicly available at github.com/jkeger/seagen or can be installed directly with pip as the python module seagen.

3.2 Uranus Giant Impacts and Convergence

We now use these tools for first creating and then simulating planets to study the convergence (or lack thereof) of giant impact simulations using 10^5 up to 10^8 SPH particles. We focus on three science-motivated questions about the giant impact that likely knocked over the planet Uranus to spin on its side: (1) How much atmosphere is ejected from the system? (2) How much rocky material is placed into orbit? (3) What is the post-impact rotation period of the planet?

Here, we repeat some of the simulations from K18 (Kegerreis et al. 2018) with $\sim 10^5$, 10^6 , 10^7 , and 10^8 particles to investigate how these higher resolutions compare with the current standard, and to demonstrate the simulation tools described in this paper. The full details of the equations of state and initial conditions are described in K18.

Fig. 5 shows comparisons of a typical impact simu-



Figure 5. Mid-collision snapshots in the early stages of the same giant impact on Uranus at the same times from simulations with the $\sim 10^5$ SPH particles (left panels) typical in the literature, up through 10^6 and 10^7 to the 10^8 (right panels) made possible with SWIFT, resolving more of the detailed evolution of both internal structure and debris. Snapshots shown are ~ 2 , 3, 4, and 7 hours after the start of the simulation. An animation of the highest resolution impact is available at icc.dur.ac.uk/giant_impacts.

lated at different resolutions, repeating the 'low angular momentum' scenario of K18's Fig. 2. Although the overall behaviour is encouragingly similar, details like the tidal stretching of the impactor's core and the distribution of the debris clearly cannot be fully resolved by the 10^5 or 10^6 particle simulations. Fig. 6 highlights some the details that can be resolved with 10^8 particles for the grazing impact of the 'high angular momentum' scenario of K18's Fig. 3.

3.2.1 Ejected Debris

In K18 we found that the majority of the atmosphere survives the impact, but that a small fraction can be fully ejected. Fig. 7 highlights the particles that will become gravit-

ationally unbound and escape from the system. The initial collision blasts away much of the outer atmosphere and some ice, some of which will escape but most remains gravitationally bound. The 10^7 and 10^8 particle runs show that a deeper shell of now-exposed particles then gets ejected during the subsequent violent oscillations as the impactor remnants fall back in and the planet slowly starts to settle. - The time at which this ejected material becomes unbound in each simulation is shown in Fig. 8. Significant mass is blasted off the planet even several hours after the initial collision in all cases. The 10^7 and 10^8 simulations closely agree that 9% of the total atmosphere mass escapes. The 10^5 and 12%, respectively. This suggests that atmospheric erosion has converged



Figure 6. A mid-collision snapshot of a grazing impact with 10^8 SPH particles – compared with the more head-on collision in Fig. 5 – coloured by their material and internal energy, showing some of the detailed evolution and mixing that can now be resolved. In the left panel, light and dark grey show the target's ice and rock material, respectively, and purple and brown show the same for the impactor. Light blue is the target's H-He atmosphere.



Figure 7. The particles that will become unbound and escape the system, highlighted in orange on a pre-impact snapshot from the same simulations with $\sim 10^5 - 10^8$ SPH particles as in Fig. 5. Only particles in a thin cross-section are shown for clarity. The colours are the same as in Fig. 6. The times at which these particles become unbound are shown in Fig. 8. The total mass lost is broadly similar in all cases, but $10^5 - 10^6$ particles fail to resolve the detailed results.

by $\sim 10^7$ particles in this case. On the positive side, although the lower resolution simulations do not show perfectly converged behaviour, for answering the practical question of how much atmosphere is lost, all simulations give a qualitatively similar answer of $\sim 10\%$.

Most studies of impact erosion use analytical or onedimensional models to estimate the ejected atmosphere given a certain ground speed from the shock induced by the impact (e.g. Inamdar & Schlichting 2016). In this case, the initial shock removes 8% of the atmosphere, then an additional 1% is lost in the subsequent sloshing. So, much like the minor resolution dependence, general conclusions about the fraction of atmosphere lost to an impact of this scale are unlikely to change. However, for more precise studies, smaller atmospheres, and perhaps other impact scenarios, this process should not be ignored.

For comparison, also shown in Fig. 8 is the mass of unbound ice. The 10^7 and 10^8 simulations again give similar

final answers, but do not show the same behaviour at earlier times. The lower resolution simulations are discrepant by more than a factor of 2. It seems plausible that this quantity is approaching convergence, but without more particles than 10^8 (or checking $10^{7.5}$), it is clearly not safe to assume this is a fully reliable result.

These quantities are summarised in Fig. 9 at 14 hours as a function of the number of particles, showing by how much each simulation differs from the highest resolution. That the eroded atmosphere appears closer to convergence than the ice is not surprising given the order-of-magnitude lower mass of ejected ice, meaning correspondingly fewer particles are involved in attempting to resolve the process – as can be interpreted by the size of the error bars.

As an example of a property that has certainly not converged, we also plot the mass of rock that is ejected into orbit in a debris disk beyond the Roche radius, where it might be available for accretion into satellites. Not only do



Figure 8. The time evolution of the mass of gravitationally unbound atmosphere (light blue) and impactor-ice (purple) material that is ejected from the system – the same particles highlighted in Fig. 7 – for the different resolution simulations.

the 10^7 and 10^8 simulations not agree, they differ by more than the lower resolutions with no semblance of convergence. The corresponding number of orbiting rock particles in each simulation is only 4, 80, 1000, and 20,000, respectively. So, especially for a messy ejection process that is widely spread out in both space and time, it is not surprising that 1000 or fewer particles are far from able to sufficiently resolve what happens. It is possible that the 10^8 simulation has already fully resolved and converged on this result, but our only means of checking this – running even higher resolution simulations – we leave for future studies where this is a targeted science result. In comparison, the masses of orbiting ice and atmosphere particles in the debris are much higher, and converge similarly to the unbound atmosphere mass.

3.2.2 Rotation Period

The rotation period of the post-impact planet is a large-scale bulk property involving a large majority of all particles, so one might expect it to have converged by fairly low particle numbers. However, as shown in Fig. 9, while the 10^7 and 10^8 simulations agree on a rotation period of 19.9 hours to within 0.5% of each other, the 10^5 and 10^6 simulations find much shorter periods of 14.7 and 17.7 hours. This is a significant change to our results in K18, when it appeared that even fairly low-impact-parameter $2 M_{\oplus}$ impactors could impart enough spin to explain the planet today. Assuming an approximately similar reduction in spin for other impacts (or more massive impactors) would be viable.

The evolution of the planet's angular momentum for each simulation is shown in Fig. 10, which, for simplicity, we sum over all particles within the Roche radius. The total angular momentum of the entire system remains the same in



Figure 9. The change with resolution of the masses of unbound atmosphere and ice, the mass of rock placed into orbit beyond the Roche radius, and the planet's rotation rate, demonstrating a range of apparent (un)convergence. Each property is normalised by the highest-resolution result to show the relative differences. The shaded regions show the 1- σ errors, some of which are too narrow to see. The rotation period appears to have converged by ~10⁷ particles, as have – with decreasing certainty – the unbound atmosphere and ice masses, while the orbiting rock mass has not.

all cases, but at higher resolutions more angular momentum is transported out to the debris disk beyond the Roche radius, leaving less in the planet. All the simulations agree during the arrival and initial merging of the impactor, but their behaviour begins to diverge as the thrown-out debris (see the middle two rows of Fig. 5) begins to fall back in to the planet, at around 3 hours after the start of the simulation.

Even though the total number of particles used to measure the planet's rotation rate is very large, the messy ejecta and mixing around the outer regions of the planet are significant enough to affect the overall system while also small enough to require high resolutions to model correctly. This is comparable to the effect seen by Hosono et al. (2017) where the mass of the post-impact disk did not converge as expected because of subtle differences in the detailed behaviour of re-impacting debris.

There will always be even smaller structures that are not properly resolved, but their ability to alter the rest of the system will also decrease, so appropriate-scale quantities should stay converged. However, properties such as small-scale turbulent mixing and the emergence of smaller structures may never converge without the addition of regularising physics such as diffusion or viscosity mechanisms (Lecoanet et al. 2016; Cullen & Dehnen 2010).

On the convergence of the rotation rate, in addition to the similar angular momenta of 10^7 and 10^8 throughout time, the rotation period encouragingly changes monotonically with higher resolution and by less with each increase. So we interpret the various quantities shown in Fig. 9 as



Figure 10. The early time evolution of the planet's angular momentum for the different resolution simulations, summed over all particles within the Roche radius of 6 R_{\oplus} . The standard- and high-resolution simulations begin to differ as the ejecta from the initial impact falls back to the planet.

demonstrating a range of behaviour from the apparently well-converged rotation rate and unbound atmosphere mass by 10^7 particles, through the possibly converged unbound ice, to the clearly un-converged orbiting rock.

4 CONCLUSIONS

We have presented a simple method for creating spherical arrangements of particles with precise densities, and the SWIFT code for hydrodynamical simulations, then used them to study giant impacts at high resolutions.

The SEA algorithm allows the quick creation of nearequilibrium, spherically symmetric initial conditions of particles (github.com/jkeger/seagen). It ensures that every particle has an SPH density within 1% of the desired value, unlike the otherwise-similarly successful methods of Raskin & Owen (2016) and Reinhardt & Stadel (2017). This mitigates the need for expensive computation that is otherwise required to produce initial conditions that are relaxed and ready for a simulation.

The open-source SWIFT code is designed to take advantage of contemporary shared/distributed-memory architectures (swift.dur.ac.uk). For planetary giant impact simulations, this has enabled a $100-1000 \times$ improvement in the number of particles that can be used, allowing the study of brand new topics that were out of reach for lower resolution simulations.

To demonstrate these tools and test the convergence of such simulations, we revisited the study of the giant impact onto the young Uranus that may explain its spin and other strange features (Kegerreis et al. 2018). We find that even large-scale results such as the rotation rate are not converged with standard-resolution simulations of 10^5 and 10^6 particles. The overall behaviour is similar in all cases, but small variations in the debris that falls back after the initial impact have a significant effect on the post-impact planet and its rotation rate, which appears to be well-converged with 10^7 and 10^8 particles, but not fewer. Similar but mildly less certain convergence is seen for the masses of atmosphere and ice that are ejected from the system, while the low mass of rock placed into orbit has not converged at all by 10^7 particles.

Increasing resolution is only one important challenge for developing more realistic simulations. We have here used a simple implementation of SPH with a focus on simply increasing the number of particles. Future studies must continue to test high resolutions with, for example, more sophisticated equations of state and improved SPH formulations with better treatment of issues such as material and density discontinuities.

We conclude that standard-resolution simulations with $<\!10^7$ SPH particles can fail to produce reliable results even for large-scale properties of planetary system. 10^7 and 10^8 particles appear to pass the threshold of resolving the major processes in a giant impact. However, different collisions and other specific simulation outputs will depend more or less strongly on the behaviour of smaller structures, with correspondingly different requirements for convergence. The highly non-linear nature of giant impacts and the combinations of short- and long-term, localised and distributed processes prevent simple predictions for how many particles will be sufficient for a given result to converge.

ACKNOWLEDGEMENTS

We thank James Willis, Josh Borrow, and all the members of the SWIFT team, and thank Lydia Heck for invaluable computational advice and support. We thank the anonymous referee for their constructive comments. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology (ICC) on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1 and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. This work was supported by INTEL through establishment of the ICC as an INTEL parallel computing centre (IPCC). JAK is supported by STFC grant ST/N50404X/1 and the ICC PhD Scholarships Fund. VRE acknowledges support from STFC grant ST/P000541/1. RJM is supported by the Royal Society. MS is supported by VENI grant 639.041.749. LFAT acknowledges support from NASA Outer Planets Research Program award NNX13AK99G.

References

Balsara D. S., 1995, Journal of Computational Physics, 121, 357
Benz W., Slattery W. L., Cameron A. G. W., 1986, Icarus, 66, 515
Borrow J., Bower R. G., Draper P. W., Gonnet P., Schaller M., 2018, CoRR, abs/1807.01341

- Ćuk M., Stewart S. T., 2012, Science, 338, 1047
- Cullen L., Dehnen W., 2010, MNRAS, 408, 669

- Dehnen W., Aly H., 2012, MNRAS, 425, 1068
- Deng H., Reinhardt C., Benitez F., Mayer L., Stadel J., 2017, preprint, (arXiv:1711.04589)
- Diehl S., Rockefeller G., Fryer C. L., Riethmiller D., Statler T. S., 2015, Publ. Astron. Soc. Australia, 32, e048
- Genda H., Fujita T., Kobayashi H., Tanaka H., Abe Y., 2015, Icarus, 262, 58
- Gonnet P., 2015, SIAM J Sci. Comput., 37, C95
- Herant M., 1994, Mem. Soc. Astron. Italiana, 65, 1013
- Hosono N., Iwasawa M., Tanikawa A., Nitadori K., Muranushi T., Makino J., 2017, Publ. Astron. Soc. Jpn., 69, 26
- Hubbard W. B., MacFarlane J. J., 1980, J. Geophys. Res., 85, 225
- Inamdar N. K., Schlichting H. E., 2016, ApJ, 817, L13
- Kegerreis J. A., et al., 2018, ApJ, 861, 52
- Kurosaki K., Inutsuka S.-i., 2019, AJ, 157, 13
- Lecoanet D., et al., 2016, MNRAS, 455, 4274
- Leopardi P., 2007, PhD thesis, Sch. Math. Stat., U. New South Wales
- Lombardi J. C., Sills A., Rasio F. A., Shapiro S. L., 1999, Journal of Computational Physics, 152, 687
- Melosh H. J., 2007, Meteoritics and Planetary Science, 42, 2079 Monaghan J. J., 1992, ARA&A, 30, 543
- Monaghan J. J., 2012, Annual Review of Fluid Mechanics, 44, 323
- Morris J. P., 1996, Publ. Astron. Soc. Aust, 13, 97
- Price D. J., 2012, Journal of Computational Physics, 231, 759
- Raskin C., Owen J. M., 2016, ApJ, 820, 102
- Reinhardt C., Stadel J., 2017, MNRAS, 467, 4252
- Saff E. B., Kuijlaars A. B. J., 1997, The Math. Int., 19, 5
- Schaller M., Gonnet P., Chalk A. B. G., Draper P. W., 2016, Proc. of PASC'16, pp 2:1–2:10
- Slattery W. L., Benz W., Cameron A. G. W., 1992, Icarus, 99, 167 Springel V., 2005, MNRAS, 364, 1105
- Springel V., 2000, Mittello, 001, 1100
- Springel V., 2010, Annual Review of Astronomy and Astrophysics, 48, 391
- Tillotson J. H., 1962, General Atomic Report GA-3216
- Wallace D. C., Sidles P. H., Danielson G. C., 1960, Journal of Applied Physics, 31, 168
- Wang J., White S. D. M., 2007, MNRAS, 380, 93
- Waples D. W., Waples J. S., 2004, Natural Resources Research, 13, 97
- Willis J. S., Schaller M., Gonnet P., Bower R. G., Draper P. W., 2018, CoRR, abs/1804.06231

Woolfson M. M., 2007, MNRAS, 376, 1173

APPENDIX A: PLANETARY PROFILES

This section details the creation of radial profiles for the model planets. The main inputs for a profile are the total mass, the number of layers and their materials, the surface pressure and temperature, and estimates for the outer radius and any internal boundary radii that we will later refine. To set each layer's material, we must define the equation of state (EoS), a conversion between temperature and internal energy e.g. the specific heat capacity and cold curve, and an expression for how heat is transferred e.g. isothermal or adiabatic.

We iterate inwards in thin spherical shells from the surface to the centre – not to be confused with the much thicker shells we define in $\S2.1.1$ to arrange simulation particles in the resulting sphere. The density at the surface is first found using the EoS with the input pressure and temperature. Assuming a constant density within this very thin shell, the mass of the shell is calculated to find the pressure at the inner shell boundary that would be required for hydrostatic equilibrium. The density and temperature that provide this pressure at the inner shell boundary are then found using the EoS and the heat transfer $(\rho-T)$ relation. This process is repeated for the next shell until reaching the centre.

The temperature and pressure are continuous across any internal layer boundaries, so this iteration continues into the core, until the input total mass has been used up. If the input radii for the outer surface and any inner boundaries are accurate, then the central shell should use up the final available mass just as its inner boundary reaches the centre. However, if any of these input radii are too large or too small, then either the mass will be used up before reaching the centre or the centre will be reached with some mass still remaining. In this case, we modify the input radii and repeat the process, until the mass discrepancy is a tiny fraction of the total mass.

For our test model of a simple Earth-mass planet in §3.1, the inputs were the Earth's mass, the Tillotson granite EoS (Tillotson 1962; Melosh 2007), and an isothermal temperature of 300 K, leading to an outer radius of 1.036 R_{\oplus} . We chose a constant specific heat capacity of 710 J K⁻¹ kg⁻¹ (Wallace et al. 1960; Waples & Waples 2004).

The resulting density (and temperature or internal energy) profile can then be used to create a set of particle initial conditions, as described in §2.1. This approach is the same for more complicated planets with multiple layers and discontinuities in material and density, such as the proto-Uranus and impactor used in §3.2 with full details in Kegerreis et al. (2018).

APPENDIX B: TILLOTSON SOUND SPEED

In addition to the pressure, density, and thermal properties of a material, the EoS is also important for determining the sound speed. In smoothed particle hydrodynamics (SPH), the sound speed is used both to control the simulation timestep – to ensure that sound waves do not travel further than the distance between neighbouring particles in one step – and as part of the artificial viscosity calculation that controls the behaviour of shocks (Price 2012).

The popular Tillotson EoS does not include an expression for the sound speed, c, but it can be derived from the partial derivative of the pressure, P, with respect to the density, ρ , at constant entropy, S:

$$c^2 = \frac{\partial P}{\partial \rho} \Big|_{S} , \qquad (B1)$$

which we can calculate from Tillotson's P, ρ , and specific internal energy u, using $du = TdS + PdV = TdS - (P/\rho^2)d\rho$.

The Tillotson pressure is separated into a condensed or cold state and an expanded and hot state (Tillotson 1962). Using the standard definitions of $\eta \equiv \rho/\rho_0$, $\mu \equiv \eta - 1$, $\nu \equiv 1/\eta - 1$, and $\omega \equiv u/(u_0\eta^2) + 1$, these two pressure formulae are

$$P_{\rm c} = \left(a + \frac{b}{\omega}\right)\rho u + A\mu + B\mu^2 \tag{B2}$$

$$P_{\rm e} = a\rho u + \left(\frac{b\rho u}{\omega} + A\mu e^{-\beta v}\right) e^{-\alpha v^2}, \tag{B3}$$

where ρ_0 , a, b, A, B, α , β , u_0 , u_{iv} , and u_{cv} are materialspecific parameters for the EoS (Melosh 2007). In the hybrid state, the pressure is a linear combination of the two:

$$P_{\rm h} = \frac{(u - u_{\rm iv}) P_{\rm e} + (u_{\rm cv} - u) P_{\rm c}}{u_{\rm cv} - u_{\rm iv}} \ . \tag{B4}$$

For SWIFT, the minimum pressure is set to 0.

Using Eqn. B1, the sound speeds for each state are

$$c_{\rm c}^2 = \frac{P_{\rm c}}{\rho} \left[1 + a + \frac{b}{\omega} \right] + \frac{b(\omega - 1)}{\omega^2} \left[2u - \frac{P_{\rm c}}{\rho} \right] + \frac{1}{\rho} \left[A + B \left(\eta^2 - 1 \right) \right]$$
(B5)

$$c_{\rm e}^2 = \frac{P_{\rm e}}{\rho} \left[1 + a + \frac{b}{\omega} e^{-\alpha \nu^2} \right] + \left\{ \frac{b\rho u}{\omega^2 \eta^2} \left[\frac{1}{u_0 \rho} \left(2u - \frac{P_{\rm e}}{\rho} \right) + \frac{2\alpha \nu \omega}{\rho_0} \right] \right. \\ \left. + \frac{A}{\rho_0} \left[1 + \frac{\mu}{\eta^2} \left(\beta + 2\alpha \nu - \eta \right) \right] e^{-\beta \nu} \right\} e^{-\alpha \nu^2} , \tag{B6}$$

and the hybrid state is the equivalent linear combination:

$$c_{\rm h}^2 = \frac{(u - u_{\rm iv}) c_{\rm e}^2 + (u_{\rm cv} - u) c_{\rm c}^2}{u_{\rm cv} - u_{\rm iv}} . \tag{B7}$$

For SWIFT, a minimum sound speed is set using the uncompressed density and bulk modulus: $\sqrt{A/\rho_0}$.

Reinhardt & Stadel (2017) did this same calculation (with slightly different notation), but their c_c^2 has a typo A instead of a in the first term and their c_e^2 has swapped the sign of $(2u - P_e/\rho)$, which would change the sound speed by ~10%.