



Universiteit  
Leiden  
The Netherlands

## Planetary giant impacts: convergence of high-resolution simulations using efficient spherical initial conditions and SWIFT

Kegerreis, J.A.; Eke, V.R.; Gonnet, P.; Korycansky, D.G.; Massey, R.J.; Schaller, M.; Teodoro, L.F.A.

### Citation

Kegerreis, J. A., Eke, V. R., Gonnet, P., Korycansky, D. G., Massey, R. J., Schaller, M., & Teodoro, L. F. A. (2019). Planetary giant impacts: convergence of high-resolution simulations using efficient spherical initial conditions and SWIFT. *Monthly Notices Of The Ras (0035-8711)*, 487(4), 5029-5040. doi:10.1093/mnras/stz1606

Version: Accepted Manuscript

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/83377>

**Note:** To cite this publication please use the final published version (if applicable).

# High-Resolution Simulations of Giant Impacts: Efficient Spherical Initial Conditions and Next-Generation Performance with SWIFT

J. A. Kegerreis<sup>1\*</sup>, V. R. Eke<sup>1</sup>, P. G. Gonnet<sup>2</sup>, D. G. Korycansky<sup>3</sup>, R. J. Massey<sup>1</sup>, M. Schaller<sup>4</sup>, L. F. A. Teodoro<sup>5</sup>

<sup>1</sup>*Institute for Computational Cosmology, Durham University, Durham, DH1 3LE, UK*

<sup>2</sup>*Google AI Perception, Google Switzerland, 8002 Zurich, Switzerland*

<sup>3</sup>*CODEP, Earth Sciences, University of California, Santa Cruz, CA, USA*

<sup>4</sup>*Leiden Observatory, Niels Bohrweg 2, 2333 CA Leiden, Netherlands*

<sup>5</sup>*BAERI/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× improvement in resolution passes a key threshold to reveal previously unseen details in the post-impact evolution of the planet’s atmosphere and ejected debris. 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 present 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 SWIFT and the particle placement code are publicly released.

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

## 1 INTRODUCTION

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. More particles means more expensive computation, so we must pursue simulation codes that take full advantage of contemporary supercomputing architectures.

In this paper, we present the simple SEA<sup>1</sup> scheme for creating optimal spherical arrangements of particles (§2) and the next-generation code SWIFT<sup>2</sup> that we have developed to run planetary impact (and cosmological) simulations (§3). We use these tools to model giant impacts onto a young Uranus with unprecedented high resolutions, using over 10<sup>8</sup> SPH particles (§4).

<sup>1</sup> The SEAGen code is publicly available at [github.com/jkeger/seagen](https://github.com/jkeger/seagen) and the python module `seagen` can be installed directly with `pip`.

<sup>2</sup> SWIFT is in open development and publicly available at [swift.dur.ac.uk](https://swift.dur.ac.uk).

\* [jacob.keggerreis@durham.ac.uk](mailto:jacob.keggerreis@durham.ac.uk)

## 2 PARTICLE PLACEMENT

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 pre-production 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, perhaps not surprisingly given its layered nature (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 (unsurprisingly) 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  $\sim 50$  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 much 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 examine 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 dis-

tribution 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. Of course, the better the initial arrangement of the particles, the less time is required for a simulation to evolve a relaxed distribution.

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 minor 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 §2.2). Reinhardt & Stadel (2017)’s method divides the sphere into equal regions that can be further subdivided, with the slight disadvantage that only restricted numbers of particles can be placed in each shell. Furthermore, some particles in each shell show SPH densities more than 5% discrepant from the desired profile density (their Fig. 4).

In the rest of this section, 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.

## 2.1 Method

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 equal-area (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  $\rho_c$ , then the thickness,  $dr$ , of a subsequent shell with density  $\rho$  is

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

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. 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 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.



**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.

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  $r_{m-w}$ , 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  $r_{m-w} > r_{1/2}$  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  $r_{m-w}$  gives too-low densities. We found that placing the particles at  $\frac{1}{2}(r_{1/2} + r_{m-w})$  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 in 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

$$A_{\text{reg}} = 4\pi/N. \quad (2)$$

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

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

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

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),

$$N_{\text{col}} = \text{round} \left[ \frac{\pi - 2\theta_{\text{cap}}}{\sqrt{A_{\text{reg}}}} \right], \quad (4)$$

and the actual initial collar height,

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

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

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

so the ideal number of regions in each collar  $i$  is

$$N'_i = \frac{A_i}{A_{\text{reg}}}. \quad (7)$$

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:


$$N_i = \text{round} [N'_i + d_i] \quad (8)$$

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

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:

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

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



**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.

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

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

$$\phi = \phi_0 + j \Delta\phi_i, \quad (12)$$


where  $\phi_0$  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,  $\phi_0$ , 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  $\phi_0$  is half the smaller of  $\Delta\phi_i$  and  $\Delta\phi_{i-1}$ . If one is odd and the other is even, then  $\phi_0$  must be half of the even one's  $\Delta\phi$ , to prevent two particles in adjacent collars from having the same  $\phi$  and being too close together.

Finally,  $\phi_0$  should be additionally offset by  $m\Delta\phi_{i-1}$ , where  $m$  is a random integer between 0 and  $N_{i-1}$ . Thus, the  $\phi_0$  rotation will be with respect to a random particle in the previous collar. This prevents successive collars with large  $N_i$  (and hence small  $\phi_0$ ) 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



**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.

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  $80 \leq N \leq 10^6$ ). 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  $b = 10a$ , and is only relevant for the innermost one or two lowest mass shells.

## 2.2 Results and Discussion

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, considering 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



**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.

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 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 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.

Of course, 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,  $v_{\text{rms}}$ , is below  $\sim 1\%$  of the escape speed,  $v_{\text{esc}} = 11.2 \text{ km s}^{-1}$ . Thanks to their precise densities, the SEA particles immediately have  $v_{\text{rms}}$  below  $0.01 v_{\text{esc}}$ , and the maximum particle speed first peaks at under  $0.04 v_{\text{esc}}$ . (‘Immediately’ here meaning the peak 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  $v_{\text{rms}} = 0.2 v_{\text{esc}}$ .

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’  $v_{\text{rms}}$  is almost identical to SEA in spite of their comparatively noisy densities. Their maximum speed is slightly higher at  $0.07 v_{\text{esc}}$ . 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 §4 – 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 ( $\sim 0.01$  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 Pluto-mass 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  $v_{\text{rms}}$ . SEA has the mild advantage that it avoids such issues in the first place, and requires similarly trivial computation to generate the particles.

The SEAGen code for quickly generating both isolated shells and full spheres of particles is publicly available at [github.com/jkeger/seagen](https://github.com/jkeger/seagen) or can be installed directly with pip as the python module `seagen`.

### 3 NEXT-GENERATION 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](https://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; Čuk & Stewart 2012), but with a  $>30\times$  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 an entirely dynamic way to achieve excellent 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 for a very good scalability of the code (Schaller et al. 2016). The gravity and hydrodynamic forces are evaluated at the same time as separate tasks. Core routines, including the direct interaction between particles, have also 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 near-perfect strong scaling even across hundreds of thousands of cores (Schaller et al. 2016; 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



**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. An animation of the highest resolution impact is available at [icc.dur.ac.uk/giant\\_impacts](http://icc.dur.ac.uk/giant_impacts).

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.

### 3.1 Planetary SPH

SWIFT has a modular structure that separates different code sections for easy 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 simply 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,<sup>3</sup> 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 required in a multi-layered planet, for example.

### 3.2 Importance and Limitations of High Resolution

Is all this effort worthwhile? SWIFT’s performance empowers us to run simulations faster and, more importantly, to use more particles than were previously feasible. We can study either larger systems with the same resolution (a com-

<sup>3</sup> SWIFT v0.7.0-2722-g628c6b27



**Figure 6.** A mid-collision snapshots 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.

mon aim in cosmology, to simulate a larger patch of the universe), or we can study a small system with higher resolution to model smaller details.

In other words, much like the size of a telescope aperture, the number of simulation particles limits the objects that can be resolved and studied. Standard works today use  $10^5$  up to  $10^6$  SPH particles which is insufficient for studying e.g. impacts onto planets with thin atmospheres or the details of ejected debris. However, attempting to simply add more particles to enable new discoveries is a serious computational challenge and, much like simply building a bigger telescope, is very (computationally) expensive.

Continuing the example of giant impacts on planets with atmospheres, most previous studies have for this reason instead used analytical approaches and 1D simulations, to estimate erosion from a range of head-on impact energies (e.g. [Inamdar & Schlichting 2016](#)). To our knowledge, the Uranus impact simulations of [Kegerreis et al. \(2018\)](#) were thus the first in three dimensions to quantify giant-impact erosion with inter-particle self-gravity, as well as the first to test a range of impact angles. This leaves much of this complex topic’s huge parameter space still to be explored, especially for lower mass atmospheres ([Liu et al. 2015](#)).

In addition to opening up new studies that were previously out of reach, the importance of improving resolution for existing topics as well was 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  $10^7$ – $10^8$ .

Of course, just because a simulation has numerically converged on a result does not make that result physically correct! SPH has several well-known difficulties, such as the interaction of multiple materials, which may not be immediately fixed by higher resolutions. That said, it is crucial that we at least converge on the answer to the (imperfect or


not) question that we ask a computer to solve, so this is an important first step.

#### 4 URANUS GIANT IMPACTS WITH $10^5$ – $10^8$ PARTICLES

We now use these tools for first creating and then simulating planets to study one of the most striking examples of the consequences of giant impacts in our solar system, Uranus, at very high resolution. Uranus spins on its side. With an obliquity of  $98^\circ$  and its major moons orbiting in the same tilted plane, the common explanation is that a giant impact sent the young planet spinning in this new direction.

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, possibly trapping the interior heat to help explain the freezing outer temperatures. We could also just about resolve the low-density atmosphere and ejected debris, finding that  $\sim 10\%$  of the atmosphere becomes unbound to escape from the system and a small amount of the impactor’s rocky core becomes mixed into the surviving outer atmosphere. [Kurosaki & Inutsuka \(2019\)](#) recently explored a different, complementary part of this 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.

Here, we repeat some of the simulations from K18 with  $\sim 10^5$ ,  $10^6$ ,  $10^7$ , and  $10^8$  particles, to demonstrate the simulation tools described in this paper and to investigate what these much higher resolutions can reveal. The full details of



**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. 7. Only particles in a thin cross-section are shown for clarity, the colours are the same as in Fig. 6. The total mass lost is similar in all cases, but  $10^5$ – $10^6$  particles fail to resolve the detailed results.


the equations of state and initial conditions are described in K18.

Fig. 5 shows a comparison of a typical impact simulated at different resolutions, repeating the ‘low angular momentum’ scenario of K18’s Fig. 2. 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. Although the overall behaviour is encouragingly similar, details like the tidal stretching of the impactor’s core and the distribution of the debris cannot be fully resolved by the  $10^5$  or  $10^6$  particle simulations.

For example, Fig. 7 highlights the particles that the impact will eject from the system. The total atmospheric erosion is consistent in all cases, but the higher resolutions reveal new details. 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 the deeper shell of now-exposed particles then gets ejected during the subsequent violent oscillations as the impactor remnants fall back in and the planet starts to settle.

In addition to the total erosion, the composition of the atmosphere is of interest both for constraining the details of collisions like this on Uranus and for understanding the chemistry and spectra of exoplanet atmospheres, many of which are likely to have suffered similar impacts (Inamdar & Schlichting 2016). Furthermore, the most common exoplanets appear most similar to our own Uranus and Neptune with a diverse range of atmospheres (Fressin et al. 2013; Pétigura et al. 2013; Lopez & Fortney 2014). Fig. 8 shows that trace amounts of the impactor’s rocky core become mixed into the outer envelope, but that  $10^5$ – $10^6$  particles are far too few to model this type of process.

Higher resolution simulations let us study finer details and confirm the convergence of our results. The  $\sim 10^7$  and  $10^8$  SPH particles that we used here are enough to reveal some of the more complex behaviour in giant impacts of this scale. Once enough particles are used to resolve and converge on the evolution of a system, then increasing the particle number further should not change the results. That said, these violent and messy planetary collisions are highly non-linear events, so increasing the resolution even further would undoubtedly uncover the intricacies of their evolution at ever smaller scales.



**Figure 8.** The masses of target and impactor materials that get mixed into Uranus’ outer envelope after the impact, for the different resolution simulations. The annotations show the corresponding number of SPH particles.

## 5 CONCLUSIONS

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

The SEA algorithm allows the quick creation of near-equilibrium, spherically symmetric initial conditions of particles ([github.com/jkeger/seagen](https://github.com/jkeger/seagen)). It ensures that every particle has an SPH density within  $\sim 1\%$  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 from scratch to take advantage of contemporary shared/distributed-memory architectures, for a significant improvement in speed and efficiency ([swift.dur.ac.uk](http://swift.dur.ac.uk)). For planetary giant

impact simulations, this has enabled a 100–1000× 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 test and demonstrate these tools, 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). Although the large-scale results from lower resolution simulations with the typical  $\sim 10^5$ – $10^6$  particles appear to have converged, the new simulations with  $10^7$  and just over  $10^8$  particles reveal new details in, for example, the erosion of the atmosphere and the composition of the mixed debris. Future studies like this will be important for our understanding of both the violent history of the planets in our solar system and the evolution of exoplanetary systems.

This is just one important issue for developing more realistic simulations, and we have here used only a simple implementation of SPH with a focus on simply increasing the number of particles. With a combination of these higher resolutions, more sophisticated equations of state, and improved SPH formulations with better treatment of issues like material and density discontinuities, many previously inaccessible topics are now becoming possible to study.

## 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. This work was supported by the Science and Technology Facilities Council (STFC) grant ST/P000541/1, and used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility ([www.dirac.ac.uk](http://www.dirac.ac.uk)). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grants ST/H008519/1 and ST/K00087X/1, STFC DiRAC Operations grant ST/K003267/1 and Durham University. 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 funded by STFC grant ST/N50404X/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  
Dehnen W., Aly H., 2012, *MNRAS*, **425**, 1068  
Diehl S., Rockefeller G., Fryer C. L., Riethmiller D., Statler T. S., 2015, *Publ. Astron. Soc. Australia*, **32**, e048  
Fressin F., et al., 2013, *ApJ*, **766**, 81  
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  
Leopardi P., 2007, PhD thesis, Sch. Math. Stat., U. New South Wales  
Liu S.-F., Hori Y., Lin D. N. C., Asphaug E., 2015, *ApJ*, **812**, 164  
Lombardi J. C., Sills A., Rasio F. A., Shapiro S. L., 1999, *Journal of Computational Physics*, **152**, 687  
Lopez E. D., Fortney J. J., 2014, *ApJ*, **792**, 1  
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  
Petigura E. A., Howard A. W., Marcy G. W., 2013, *Proc. Nat. Aca. Sci. USA*, **110**, 19273  
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., 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

## 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 §2.1.1 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, 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  $710 \text{ J K}^{-1} \text{ kg}^{-1}$  (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. 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 §4 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,  $\rho$ , at constant entropy,  $S$ :

$$c^2 = \left. \frac{\partial P}{\partial \rho} \right|_S, \quad (\text{B1})$$

which we can calculate from Tillotson's  $P$ ,  $\rho$ , 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

$$P_c = \left( a + \frac{b}{\omega} \right) \rho u + A\mu + B\mu^2 \quad (\text{B2})$$

$$P_e = a\rho u + \left( \frac{b\rho u}{\omega} + A\mu e^{-\beta\nu} \right) e^{-\alpha\nu^2}, \quad (\text{B3})$$

where  $\rho_0$ ,  $a$ ,  $b$ ,  $A$ ,  $B$ ,  $\alpha$ ,  $\beta$ ,  $u_0$ ,  $u_{iv}$ , and  $u_{cv}$  are material-specific parameters for the EoS (Melosh 2007). In the hybrid state, the pressure is a linear combination of the two:

$$P_h = \frac{(u - u_{iv})P_e + (u_{cv} - u)P_c}{u_{cv} - u_{iv}}. \quad (\text{B4})$$

For SWIFT, the minimum pressure is set to 0.

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

$$c_c^2 = \frac{P_c}{\rho} \left[ 1 + a + \frac{b}{\omega} \right] + \frac{b(\omega - 1)}{\omega^2} \left[ 2u - \frac{P_c}{\rho} \right] + \frac{1}{\rho} \left[ A + B(\eta^2 - 1) \right] \quad (\text{B5})$$

$$c_e^2 = \frac{P_e}{\rho} \left[ 1 + a + \frac{b}{\omega} e^{-\alpha\nu^2} \right] + \left\{ \frac{b\rho u}{\omega^2\eta^2} \left[ \frac{1}{u_0\rho} \left( 2u - \frac{P_e}{\rho} \right) + \frac{2\alpha\nu\omega}{\rho_0} \right] + \frac{A}{\rho_0} \left[ 1 + \frac{\mu}{\eta^2} (\beta + 2\alpha\nu - \eta) \right] e^{-\beta\nu} \right\} e^{-\alpha\nu^2}, \quad (\text{B6})$$

and the hybrid state is the equivalent linear combination:

$$c_h^2 = \frac{(u - u_{iv})c_e^2 + (u_{cv} - u)c_c^2}{u_{cv} - u_{iv}}. \quad (\text{B7})$$

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  $\sim 10\%$ .

This paper has been typeset from a  $\text{\TeX}/\text{\LaTeX}$  file prepared by the author.