Icarus 212 (2011) 719–735

Contents lists available at ScienceDirect

Icarus

journal homepage: www.elsevier.com/locate/icarus

A numerical model of cohesion in planetary rings

Randall P. Perrine [a,*], Derek C. Richardson [a], Daniel J. Scheeres [b]

[a] Department of Astronomy, University of Maryland, College Park, MD 20740-2421, United States
[b] Department of Aerospace Engineering Sciences, Colorado Center for Astrodynamics Research, University of Colorado at Boulder, Boulder, CO 80309-0429, United States

ARTICLE INFO

Article history:
Received 15 June 2010
Revised 21 January 2011
Accepted 25 January 2011
Available online 1 February 2011

Keywords:
Planetary rings
Saturn, Rings
Rotational dynamics
Tides, Solid body
Ices, Mechanical properties
Collisional physics

ABSTRACT

We present a numerical method that incorporates particle sticking in simulations using the *N*-body code `pkdgrav` to study motions in a local rotating frame, such as a patch of a planetary ring. Particles stick to form non-deformable but breakable aggregates that obey the (Eulerian) equations of rigid-body motion. Applications include local simulations of planetary ring dynamics and planet formation, which typically feature hundreds of thousands or more colliding bodies. Bonding and breaking thresholds are tunable parameters that can approximately mimic, for example, van der Waals forces or interlocking of surface frost layers. The bonding and breaking model does not incorporate a rigorous treatment of internal fracture; rather the method serves as motivation for first-order investigation of how semi-rigid bonding affects the evolution of particle assemblies in high-density environments.

We apply the method to Saturn's A ring, for which laboratory experiments suggest that interpenetration of thin, frost-coated surface layers may lead to weak cohesive bonding. These experiments show that frost-coated icy bodies can bond at the low impact speeds characteristic of the rings. Our investigation is further motivated by recent simulations that suggest a very low coefficient of restitution is needed to explain the amplitude of the azimuthal brightness asymmetry in Saturn's A ring, and the hypothesis that fine structure in Saturn's B ring may in part be caused by large-scale cohesion.

This work presents the full implementation of our model in `pkdgrav`, as well as results from initial tests with a limited set of parameters explored. We find a combination of parameters that yields aggregate size distribution and maximum radius values in agreement with Voyager data for ring particles in Saturn's outer A ring. We also find that the bonding and breaking parameters define two strength regimes in which fragmentation is dominated either by collisions or other stresses, such as tides. We conclude our study with a discussion of future applications of and refinements to our model.

1. Introduction

1.1. Motivation

The small particles comprising the rings of Saturn are composed almost entirely of water ice (see Cuzzi et al. (2009) for a recent review). Laboratory measurements of collisions between frost-covered icy spheres (Hatzes et al., 1991; Bridges et al., 1996; Supulver et al., 1997; Bridges et al., 2001) have shown those bodies to adhere at very low impact speed. In particular, Hatzes et al. (1991) showed evidence for a critical impact speed below which every impact results in cohesion (~0.3 mm/s); this speed happens to be on the order of the escape speed (for a test particle on the surface) of an icy sphere 1 m in radius. Their analysis explains this cohesion with a "Velcro" model—the rough texture of the frost layers providing an interlocking structure for colliding particles.

Interestingly, analysis of recent observations of the rings of Saturn implies the need for such a bonding mechanism. Porco et al. (2008) show that the contrast in the azimuthal brightness asymmetry of Saturn's A rings (e.g., Salo et al., 2004; French et al., 2007; Porco et al., 2008) can be modeled using extremely dissipative particle interactions, such as would arise from bonding. And, until recently, radial banded structures in Saturn's B ring eluded explanation. Tremaine (2003) suggests that these zones occupy a region of dynamical phase-space that is "shear-free"—zones with zero Keplerian shear—and thus subject to a solid–liquid phase transition. The suggestion here is that these zones may consist of large-scale transiently bonded particles that orbit Saturn as a solid. While these banded structures have been attributed recently to viscous overstability in the dense rings (Spitale and Porco, 2010), transient cohesive bonding may still play some role in the dynamics and observable properties of these regions.

Studying the full effect of such bonding mechanisms in dense environments, specifically Saturn's dense A and B rings, requires detailed numerical modeling. Such systems involve a complex convergence of phenomena, including interparticle self-gravity,

* Corresponding author. Fax: +1 301 314 9067.
E-mail address: perrine@astro.umd.edu (R.P. Perrine).

0019-1035/$ - see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.icarus.2011.01.029

planetary tides, and interparticle collisions. It is beyond the capability of current analytical models to further incorporate dynamical bonding and fragmentation. Some groups have modeled the effect of cohesion between a small number of bodies (e.g., Spahn et al. (2004) and Albers and Spahn (2006) using a viscoelastic model), but none have attempted to self-consistently model the global effects that emerge when tens of thousands (or even millions) of icy bodies interact over many orbital times.

Local simulations are valuable tools for these studies: by restricting the computational volume to a small region of interest (which we call a "patch"), realistic surface densities and particle size distributions can be modeled. Specifically, a full-ring simulation of Saturn's A and B rings with accurate surface densities and spherical 1 m radius bodies would require roughly 10^{15} particles, which is far beyond current computational possibility. But a representative patch may require only 10^{5-9} bodies, depending on the specific problem being explored.

Local simulations can be used to study planetary rings, debris disks around stars, or any other non-inertial system orbiting a large body. Such systems use Hill's linearized equations of motion (Hill, 1878; Wisdom and Tremaine, 1988):

$$\begin{aligned} \ddot{x} &= \mathscr{F}_x + 3\Omega^2 x + 2\Omega\dot{y}, \\ \ddot{y} &= \mathscr{F}_y - 2\Omega\dot{x}, \\ \ddot{z} &= \mathscr{F}_z - \Omega^2 z \end{aligned} \tag{1}$$

with $\mathscr{F}$ being the acceleration due to particle self-gravity, Ω the Keplerian orbital frequency of the system (which equals $\sqrt{GM/a^3}$, with G the gravitational constant, M the mass of the central body, and a the distance to that body), x, y, and z the coordinates of the particle in the local coordinate system (whose origin is located at the center of the patch), and derivatives with respect to time. See Fig. 1.

The azimuthal and radial extents of the patch are small compared to its orbital distance from the planet, but large compared to the radial mean free path of the particles inside it. Periodic boundary conditions are employed in x and y to ensure that the total number of particles in the rectangular volume containing the patch remains constant. Particles exiting one side of the volume (in x or y) reappear on the other side with the same properties (mass, spin, random component of velocity, etc.) as the exiting body. We refer to this process as "wrapping" the particle. When a particle crosses the radial boundary, Keplerian shear requires that its azimuthal velocity be adjusted by $\frac{3}{2}\Omega L_x$ (where L_x is the radial dimension of the patch). This ensures a smooth velocity transition across the patch boundary. In addition, the particles in the patch are replicated into surrounding patches; these provide gravitational and collisional boundary conditions. See Fig. 1.

Many groups employ this local "patch" model to study planetary rings (recent examples include, Lewis and Stewart, 2000; Karjalainen and Salo, 2004; Porco et al., 2008), though none of those groups employ interparticle bonding in their models. Thus our model appears to be unique, as it incorporates rigid bonding, using user-defined merging and fragmentation criteria, into a local N-body simulation of a self-gravitating, collisionally evolving system.

This model is also of interest in the planet formation field, as it will improve upon the utility of existing methods. For example, Barnes et al. (2009) includes a mechanism for merging N-body particles and growing planetesimals; however, those mergers merely replace the colliding particles with a single spherical particle with the same total mass and angular momentum, losing critical information, such as shape and spin, as the planetesimals evolve. Further, that model has no fragmentation mechanism. Spahn et al. (2004) provides an analytical approach to collisional merging and fragmentation in the context of planet formation, but relies on replacing particles with spherical collision remnants. Our new method allows for studies into more realistic shape and spin distributions via fractal growth, and accounts for (somewhat simplistic) breakage due to collisions and external forces.

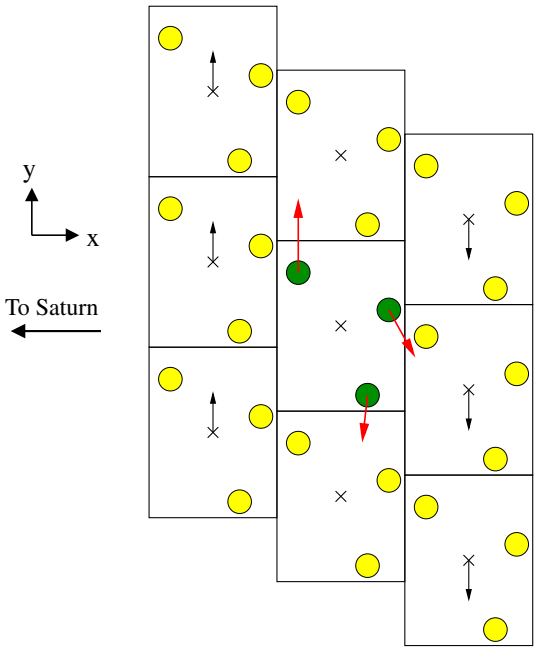

Fig. 1. A schematic diagram of a local sliding-patch model with shearing periodic boundary conditions. The three green (or dark, if viewed in gray scale) particles in the center box are the simulated particles; yellow (light) particles are in replicated patches that provide boundary conditions. The x coordinate is the radial direction, with Saturn located far in the $-x$ direction; y is the azimuthal direction, and the entire patch orbits Saturn in the $+y$ direction. z points out of the page, forming a right-handed coordinate system. Note that while the patch is periodic in x and y, it is unbounded in z. The simulation is carried out in the orbital frame of the center of the patch, so particles to the left shear upward; on the right, they shear downward. The replicated patches similarly shear (or "slide") past the center box in Keplerian fashion; each bold X marks the center of each patch, with the bulk motion of each patch indicated by black arrows. (This is the origin of the term "sliding patch.") Particles crossing boundaries reappear on the opposite side, with positions and velocities adjusted for shear; each particle in this example will soon cross a boundary, with current velocities indicated by attached red arrows. The azimuthal velocity of the particle making a radial boundary crossing will be adjusted for shear by $+\frac{3}{2}\Omega L_x \dot{y}$ (see Section 1.1). Note that while only one ring of replicated patches is shown, we typically use three rings of replicas, in order to provide a smooth background. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Our model may not yet include all of the complex physics relevant to fractal growth and granular mechanics, but it remains more sophisticated than previous methods. In applying our model to the rings of Saturn, we follow the precedent of the planet formation field, in which basic models are first applied to the problem to gain first-order understanding before investing more effort into ever more detailed models with ever more accurate physics.

1.2. Terminology

Following the nomenclature of Richardson et al. (2009), we refer to an "aggregate" of material as a collection of particles that are joined together via inflexible and incompressible bonds of user-defined strength. For example, this body can, depending on the strength of the bonds, resist self-gravitational reshaping (e.g., remain highly non-spherical). The constituents of the aggregate travel through space as one unit (i.e., the particles' positions remain fixed with respect to one another), and rotate collectively in response to external torques. Aggregates can experience fractal growth through collisional accretion with other bodies—free

particles, or other aggregates—and can fragment and lose particles through collisional events or stress fragmentation (when stresses exceed bond strengths).

Other researchers (e.g., Karjalainen and Salo, 2004; Karjalainen, 2007) use the term "aggregate" to refer to a collection of particles bound together by gravity alone. We refer to these structures as "gravitational aggregates." Unless indicated otherwise, "aggregate" in this paper refers to a bonded aggregate, though both types of structures appear in our simulations (cf. Section 3.2).

An important type of gravitational aggregate that is prevalent in dense ring simulations is a structure usually referred to as a "gravity wake," or simply a "wake" (see Salo, 1992, for a full discussion; for more theory, see, for example, Daisaka and Ida, 1999). While these wakes have never been imaged directly (due to their currently unresolvable sub-hundred-meter size), they are predicted by many models (e.g., Salo, 1992; Colwell et al., 2007; Porco et al., 2008; Robbins et al., 2010). In the models, they are distinguished from other types of gravitational aggregates by stability, morphology, and orientation, and form when the gravitational collapse due to particle self-gravity nearly balances the disruptive influence of Keplerian shear. Given sufficient surface density and distance from the central body, these wakes form readily, yet are disrupted on orbital times, and are thus highly transient structures. They take the form of elongated and densely packed gravitational aggregates, with a separation estimated by twice the critical wavelength of the region: $\lambda_{crit} = 4\pi^2 G\Sigma/\kappa^2$ (with Σ the surface mass density, and κ the local epicyclic frequency—which in this case is the Keplerian orbital frequency Ω). Due to their formation mechanism, wakes have a characteristic pitch angle (orientation angle with respect to the orbital direction) of about 20–25°, depending on local conditions.

Researchers (e.g., Salo et al., 2004; Porco et al., 2008) have invoked the presence of these wakes to explain the azimuthal brightness asymmetry in Saturn's A ring. Wakes are a dominant structure in our simulations (see Section 3.2).

1.3. `pkdgrav`

We built our aggregation model into `pkdgrav`, an *N*-body code originally designed for cosmological simulations (Stadel, 2001), which was adapted to include particle collisions (Richardson et al., 2000) for the purpose of studying, among other things, planetary rings (e.g., Porco et al., 2008).

`Pkdgrav` uses a parallelized tree code to reduce the computational cost of summing up gravitational force contributions between particles. A second-order leapfrog scheme is used to integrate the equations of motion by alternatively updating particle positions and velocities; velocity updates are performed during the "kick" phase, and position updates during the "drift" phase. (Velocities are held constant during the drift, and positions are held constant during the kick.) Thus the code proceeds forward in time as a sequence of alternating kicks and drifts.

Collisions among particles are predicted and resolved during the drift phase by treating particle trajectories as linear and predicting when intersections occur (for more discussion, see, e.g., Richardson et al., 2009).

As discussed in Quinn et al. (2010), naïvely applying the leapfrog method to integrate Hill's equations generally breaks the symplectic nature of the integrator. Quinn et al. (2010) presents a symplectic leapfrog method for the rotating patch frame that we have implemented into our code. One complication of the method is the introduction of a new canonical variable, the momentum P_y, that needs to be tracked for each particle. Impulsive events, like collisions, require updates to P_y during the drift phase. We adapted Eqs. (33) and (34) of Quinn et al. (2010) to find P_y during the drift:

$$P_y = \dot{y} + 2\Omega(x_{event} + \dot{x}(\Delta t/2 - t_{event})) \quad (2)$$

where $\dot{x}$ and $\dot{y}$ are the (constant) velocities during the drift, Ω is the orbital frequency of the rotating coordinate system, x_{event} is the (drifted) radial position of the body at the time of the event, Δt is the full timestep, and t_{event} is the elapsed time since the start of the drift phase. Usually, the "event" is a collision, though other events that require recalculation of P_y do occur (e.g., stress fragmentation).[1] It should be noted that both dissipative collisions and the non-momentum-conserving nature of the tree-based method also "break" the symplecticity of the leapfrog integrator—but the former is a controlled energy loss, and the latter is minimized in the patch model because forces rapidly homogenize with distance in the flattened geometry of planetary rings.

The remainder of this paper is divided as follows. Section 2 describes the numerical method in detail. Section 3 describes tests we performed that demonstrate the code is working correctly. Section 4 provides a summary and ideas for future work. Appendix A gives the complete derivation of the equations of motion of rigid bodies in the rotating patch frame. Appendix B presents a geometrical calculation for the contact area between rigid particles general enough for any particular bonding model. Lastly, Appendix C presents an expression and rationale for our choice of bonding parameters by relating the impact pressure between colliding bodies to the impact speed.

2. Aggregate model

2.1. Previous work

Richardson (1995) incorporated a basic version of the aggregate model (for the inertial frame) into an earlier *N*-body code. Richardson et al. (2009) incorporated the model into `pkdgrav`, on which the present work is based. We review the essential details here.

Colliding particles can stick on contact to form an aggregate—or add to an existing aggregate—if the impact speed is below a user-defined threshold called the "merge limit" (which can be chosen by the user to be either a fixed value, or a value proportional to the mutual escape speed of the colliding objects). Fractal growth proceeds as more free particles—or other aggregates—collide and stick in this way. In our model, particles have no knowledge of which particle(s) they are "stuck" to—all a particle knows is which aggregate it belongs to. Aggregate properties are calculated based on the positions, velocities, spins, etc., of the particles included in that body alone.

To counter growth, an aggregate must be able to fragment, and we have implemented two breaking mechanisms that can cause a bonded aggregate to lose particles. First, a particle that impacts an aggregate at a speed above a (different) user-specified threshold causes the impacted particle to become liberated from the aggregate. This threshold is the "fragmentation limit" (which can also be either an absolute speed or a value scaled by the mutual escape speed of the colliding bodies). If the impacting particle is itself part of an aggregate, it may be liberated as well. Liberated particles will likely immediately strike other particles in their respective former

[1] For the interested reader: P_y is defined and calculated for all particles at the beginning of the kick phase, before any particle's velocities are updated. Further, the drift phase adjusts a particle's position, using the adjusted velocity. Therefore, in order to properly compute P_y at any arbitrary time during the drift phase, we must both backtrack the particle's position to the start of the drift, and undo the velocity change that occurred during the kick phase. Thus Eq. (2) depends on both the elapsed time during the drift (to undo the drift) and the timestep (to undo the kick). Note that at the exact middle of the drift, when $\Delta t/2 = t_{event}$, the expression simplifies to $P_y = \dot{y} + 2\Omega(x_{event})$; this is because Eq. (2) is based on the cross term of the Hamiltonian (that is the basis of the symplectic method) as measured from the middle of the step (see the discussion in Quinn et al. (2010), preceding Eq. (29)).

aggregates, which might also become liberated (if the impact speeds are still high enough), causing a fragmentation cascade.

Second, each aggregate is assigned a user-defined strength,[2] or "stress limit," in the normal (tensile) or tangential (shear) directions, or both. The aggregate experiences no strain as the stress increases: it remains perfectly rigid until the strength is exceeded. The failure test is performed at each step immediately following the gravity calculation: each bonded aggregate is checked in turn to see if any constituent particles are experiencing a differential acceleration relative to the aggregate center of mass that exceeds the strength. This stress can arise from such factors as rapid rotation, or gravitational tidal effects from the planet or other particles, and is offset by the effect of gravitational self-compression (since all interparticle forces are computed explicitly—even those arising between particles within an aggregate). Any particle whose acceleration exceeds the stress limit is liberated from the aggregate and becomes a free particle again at its current position and velocity.

For these calculations, the strength, in units of pressure, is converted to a maximum acceleration by multiplying by the cross-sectional area and dividing by the mass of the particle under consideration. We recognize that the cross-sectional area of the cohesive particles is a gross over-approximation to the true contact area of most particle bonding. But any bonding model (whether it is based on deformable particles, or some means of interpenetration of the surfaces of the bodies) will have a drastically different contact area. Our method makes no attempt to select one model over another, in order to remain a general and simple method to explore the effects of cohesion—no matter the means of sticking under study. Thus it is up to the user of the code, and our readers, to translate the strengths used as parameters within `pkdgrav` into strengths applicable to the bonding method under study. This is achieved by scaling the strength parameter by the ratio of the cross-sectional area to the true contact area (see Appendix B).

Note that there is no internal consistency check to verify that the user's choice of fragmentation limit matches physically to the choice to bond strength. Thus the user can chose to model aggregates with infinite strength (such that tides never disrupt an aggregate) while simultaneously allowing those aggregates to fragment at the slightest of impacts, or aggregates that are highly resistant to impact fragmentation yet are easily tidally disrupted. This is a useful feature, as it can be illuminating to isolate the fragmentation mechanisms to study their behaviors. So it is up to the user to match these two parameters physically, if desired. In order to assist with understanding, we have derived a naïve order-of-magnitude relationship between the stress limit and the velocity-based fragmentation limit to use as a general guideline (see Appendix C). This relationship includes the generic parameter β (defined in Appendix B) that scales the strength by considering the varying contact areas of different bonding methods.

Also note that we do not keep track of explicit bonding networks within an aggregate, such as those discussed in Lois and Carlson (2007); whether a particle joins or leaves an aggregate is based entirely on the properties of that particle alone and the aggregate as a whole, not on local conditions within the aggregate. For instance, we cannot track failures or cracks in this model—all bonds obey the same strength law.

Additionally, collisional fragmentation is limited to removing individual particles from an aggregate; once liberated, each particle is itself an indestructible sphere, yet each particle can continue to fragment the aggregate in a cascade of collisions. Thus it is impossible, for example, for an impacted aggregate to break into two or more large pieces in a single event. This could result in an over-abundance of small particles in equilibrium, as those particles ought to be combined as larger collisional fragments. However, our fragments can spontaneously recombine following the cascade—particularly if those particles emerge from the impact event on similar trajectories—to form a polydisperse distribution of fragments. Thus our cohesion model is only an approximation to the behavior of real cohesive materials, but we believe we are capturing the essential elements while keeping the computations tractable.

However, the model's generality remains a strength, as it captures the basic behavior of bonding without limiting the range of cohesion mechanisms that it can mimic. For example, our basic model could approximate soft deformable spheres, such as those used in Johnson et al. (1971), or overlapping frost layers, such as in Hatzes et al. (1991), by merely a careful change to the bonding parameters.

Having discussed the bonding aspects of our model (how aggregates are created and destroyed), we next move on to a summary of the dynamics of our aggregates. Aggregates obey Euler's equations of rigid-body rotation:

$$\begin{aligned} I_1\dot{\omega}_1 - \omega_2\omega_3(I_2 - I_3) &= N_1 \\ I_2\dot{\omega}_2 - \omega_3\omega_1(I_3 - I_1) &= N_2 \\ I_3\dot{\omega}_3 - \omega_1\omega_2(I_1 - I_2) &= N_3 \end{aligned} \tag{3}$$

where I_k are the principal moments of inertia of the body, ω_k are the spin components in the body frame, $\dot{\omega}_k$ are the time derivatives of those components, and N_k are the net torque components in the body frame. The evolution of the orientations of an aggregate's principal axes is given by:

$$\begin{aligned} \dot{\hat{\mathbf{p}}}_1 &= \omega_3\hat{\mathbf{p}}_2 - \omega_2\hat{\mathbf{p}}_3 \\ \dot{\hat{\mathbf{p}}}_2 &= \omega_1\hat{\mathbf{p}}_3 - \omega_3\hat{\mathbf{p}}_1 \\ \dot{\hat{\mathbf{p}}}_3 &= \omega_2\hat{\mathbf{p}}_1 - \omega_1\hat{\mathbf{p}}_2 \end{aligned} \tag{4}$$

where $\hat{\mathbf{p}}_i$ denote the principal axes. Note that these are simplified for use in an aggregate's body frame, in which the inertia tensor is diagonalized. This further simplifies matters as the body frame does not include the rotational accelerations that are present in the rotating patch frame.

Torques in Euler's equations have the form $(\mathbf{r}_i - \mathbf{r}_a) \times (\ddot{\mathbf{r}}_i - \ddot{\mathbf{r}}_a)$—i.e., they depend on a constituent particle's position ($\mathbf{r}_i$) relative to the aggregate center of mass ($\mathbf{r}_a$), and the total relative force per unit mass acting on the particle (including tides from the planet, etc.). We use a fifth-order (time-adaptive) Runge–Kutta integrator to advance the spin vectors and orientations of the aggregates during the drift phase according to Eqs. (3) and (4). The aggregate centers of mass are advanced according to Eq. (1) in the usual way while particles in the aggregates are constrained to obey Euler's equations. Collisions involving aggregates that result in bouncing (i.e., too fast for sticking but too slow for breaking) are treated using non-central impact equations (with no surface friction); see Richardson et al. (2009) for exact equations. Note that particles inside aggregates do not move relative to one another, which saves on collision searches (considerably so for large aggregates that would otherwise be computationally expensive rubble piles).

Because of the extra rotation component of aggregates, collision prediction involving them is more complex—even though the solution still just consists of predicting the intersection of spheres. (The complication is that the spheres are no longer on straight-line trajectories, due to the aggregate rotation.) Formally, always within the context of second-order leapfrog (for which translational velocities are held constant during the drift phase), collision prediction for rotating aggregates requires solving a quartic equation. For sufficiently small timesteps, and reasonable rotation rates, the

[2] The strength can be made size-dependent according to $S \propto r^{\alpha}$, where r is the radius of the aggregate's minimally enclosing sphere, and α is a user-defined constant.

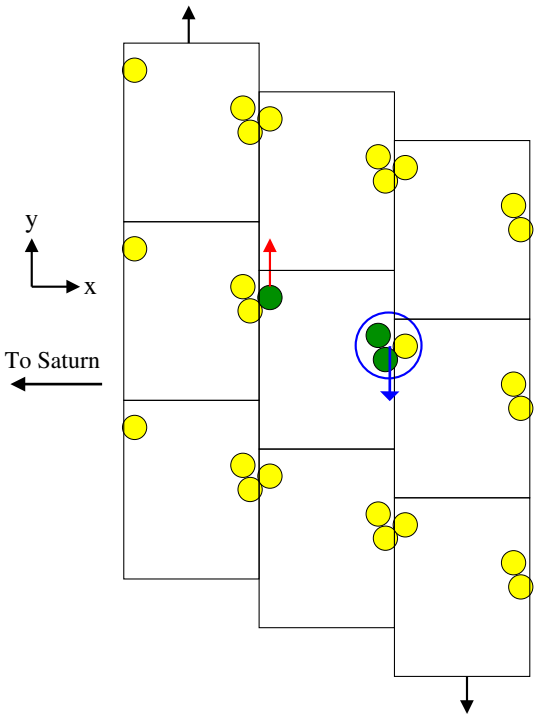

Fig. 2. A schematic diagram similar to Fig. 1 illustrating the considerations involved with handling aggregates in a periodic patch. There are three particles in this system, and they are joined into a single aggregate. The aggregate is encircled, and the center of mass (COM) of this aggregate would lie at the approximate center of the circle (as drawn). The green (dark) particles represent the positions of the particles with the periodic boundary conditions taken into account (labeled as position #1 in Section 2.3.2), and the yellow (light) particles are their replicas in eight adjacent patches. As the aggregate lies across a boundary, the circle encloses both yellow and green particles. If the model were to compute the location of this aggregate's COM from the positions of the green particles, the result would be incorrect. A similar scenario exists with particle velocities: as in Fig. 1, Keplerian shear carries particles on the right downward, and particles on the left upward. Thus the green particle on the left of the patch has a velocity in the $+y$ direction (upward red arrow). However, the aggregate's COM lies on the right of the patch, and therefore is carried in the $-y$ direction by shear (downward blue arrow). To calculate an aggregate's COM velocity, the model sums over the velocities of its constituent particles; if it naïvely used green particle velocities, it would include the red arrow, which would produce an incorrect result. Thus, we must store the "unwrapped" position and velocity of any such yellow particles in aggregates (labeled position and velocity #2) and use that data to calculate the state of the COM (cf. Section 2.3.2). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

quartic can safely be approximated as a quadratic (the user has the option of solving the full quartic, at a performance penalty).

However, even with the full quartic treatment, particle overlaps occur, because the collision prediction only extrapolates the aggregate motion (i.e., without solving Euler's equations explicitly). See Section 2.4 for a detailed discussion of particle overlaps.

2.2. Modifications: overview

Two essential issues required direct attention while incorporating the rigid aggregate model of Richardson et al. (2009) into the local rotating frame model, which we outline here. Details are in Section 2.3.

The first is that with the introduction of periodic boundary conditions, portions of a given aggregate can appear on opposite sides of the patch simultaneously (see Fig. 2)—complicating, for example, collision outcome handling, and calculation of the properties of the aggregate's center of mass (COM). To address this, we store the positions and velocities of each particle in an aggregate with respect to more than one reference point: when an aggregate lies across a patch boundary (i.e., the circle in Fig. 2 intersects with a boundary), we store not only the standard in-patch position (shown in green in Fig. 2), but also the position of the particle that is guaranteed to be near its COM—a position that may be located outside of the patch (a case represented by the yellow particle in the circle). We call this process "unwrapping" the particle, to contrast the in-patch "wrapped" particle. This greatly simplifies computation of the aggregate's COM properties (see Section 2.3.2).

The second issue is that the aggregate dynamics must be transformed properly to the rotating frame. This involves special considerations for aggregate rotation, orientation, external forcing from planetary torques, and collision detection. For example, in an inertial frame, in the absence of external influence, an aggregate maintains the orientation of its angular momentum vector with respect to the coordinate axes. But if viewed in a local rotating frame, this same aggregate's angular momentum vector would appear to precess with a frequency equal to the orbital frequency of the patch (Ω). This in turn has subtle effects on such processes as collision detection and resolution, and requires rederiving the equations of motion for an aggregate in a rotating frame (see Appendix A).

2.3. Modifications: details

This section describes explicitly all of the modifications and additions we made to the numerical method in placing the rigid aggregate model into the rotating local frame. This section can safely be skipped if the reader is not interested in this level of detail.

2.3.1. Three coordinate systems

We begin by presenting the three sets of Cartesian coordinate systems ("frames") in this model. Then we discuss their function, and finally how we transform between them.

The "body" frame's origin is at the COM of the aggregate, and its axes are aligned with the principal axes of the body. In this frame, the inertia tensor is diagonalized, and the Euler equations of rigid-body motion are easily solved numerically. As the aggregate rotates, the orientation of these axes changes in an external frame, but not in the body frame (i.e., the body does not appear to move in this frame).

The "space" frame is an inertial frame whose axes are stationary. The origin's location in this frame is arbitrary. In this work, this frame is used chiefly when storing the spins of free particles, and as an intermediate step in transforming from the body frame to the patch frame.

The "patch" frame is the frame in which we work most often, as the sliding-patch model is based in this non-inertial frame (see Fig. 1). Its axes are initially aligned with the space frame at $t = 0$; but as the patch orbits, it rotates about its $\hat{z}$ axis at Ω, the orbital frequency of the patch.[3] Translational motions for all bodies are calculated in the patch frame, using Hill's equations of motion (Eq. (1)), which include the coriolis and centrifugal accelerations.

To transform from the patch frame to the space frame, one must rotate the coordinate axes (about $\hat{z}$) by an angle Ωt, where t is the time since the simulation began. Also, any aggregates or particles must add $\Omega\hat{z}$ to their spin vector.

To transform from the body frame to the patch frame, multiply any body-frame vector by Λ, the matrix composed of the eigenvectors of the inertia tensor (see Richardson et al., 2009), where the eigenvectors are the orientations of the principal axes of the body as seen in the patch frame. This is equivalent to performing two

[3] The origin of the patch coordinate system does not coincide with the axis of rotation; the distance between the origin and the center of rotation is the patch's orbital semimajor axis, a.

matrix multiplications: one to transform into the space frame, and another to further transform into the patch frame. As such, both rotations are performed in one operation. Note that to transform spins from the body to the patch frame, one must finally subtract $\Omega\hat{z}$ from all spin vectors.

Conversely, in order to go from patch to body frames, simply multiply vectors by A^{-1}. To properly handle spin, first add $\Omega\hat{z}$ to the spin (to put the spin into the space frame), and then multiply by A^{-1}.

2.3.2. Aggregates and periodic boundaries

As mentioned in Section 2.2, there is a bookkeeping issue when aggregations of particles interact with periodic boundary conditions: it is no longer obvious on which side of the patch the center of mass (COM) of any aggregate lies. For an aggregate overlapping a boundary, a naïve weighted average of particle positions would place the COM far from its true COM location (see Fig. 2). Also, when aggregates overlap a radial boundary, they obey shear imposed by the sliding patch (i.e., parts of the aggregate will move upward, and the rest downward), which similarly complicates calculation of the COM's velocity.

Rather than devise a computationally expensive scheme to detect and compensate for periodic boundary offsets multiple times during every timestep (e.g., for multiple collisions), we simply track the aggregate's constituent particles more completely. If it is joined to an aggregate, we utilize two sets of coordinates to store the position of a single particle:

1. Wrapped patch position.
2. Unwrapped patch position.

Position #1 is standard in the sliding-patch model: the position of the particle in patch coordinates that has been wrapped as necessary to keep it inside the patch. This position is not necessarily on the same side of the patch as the aggregate's COM (that is, the distance to the COM is not the magnitude of the difference in position vectors). Position #1 is used for gravity calculation and collision prediction. These are the green particles in Fig. 2.

Position #2 is in patch coordinates, but is not wrapped as is usual in a patch model—it is allowed to be outside the boundaries of the patch. This position is guaranteed to be near the aggregate COM (that is, subtracting the position vectors gives the correct displacement from the COM). It is used to calculate the aggregate's COM position, spin vector, angular momentum, inertia tensor, and torque vector. In Fig. 2, this is the position of the yellow particle in the circle.

We also store the velocities of the particles in two ways, due to the differential shear across the patch. The two types of velocity are:

1. Wrapped patch velocity (shear matches position #1).
2. Unwrapped patch velocity (shear matches position #2).

Each velocity is measured in patch coordinates. Velocity #1 is standard in the patch model, with particle shear applied as usual. This velocity is used for collision detection between particles. Note that collision detection between aggregates is complicated by the rotating frame (see Section 2.3.5).

Velocity #2 is an "unwrapped" version of #1—that is, particles in aggregates that extend beyond the edge of the patch do not readjust their azimuthal velocities due to the Keplerian shear across the boundary. This velocity is used to determine aggregate velocity, spin, and angular momentum—free of any boundary-condition-induced complications. (Fig. 2 refers to this issue specifically.) Note that velocities #1 and #2 are identical except for their azimuthal components.

When an aggregate approaches a patch boundary, we must take care to apply boundary conditions properly to the aggregate's constituent particles:

A particle's position #1 is wrapped whenever that particle crosses a boundary, to guarantee that it remains within the patch. Additionally, as in the standard sliding-patch model, shear is applied to velocity #1 when a radial boundary is crossed.

In contrast, a particle's position #2 is only wrapped when the COM of its aggregate crosses a boundary. In fact, all particles in that aggregate have identical offsets applied their position #2 simultaneously, to maintain the relative positions of the aggregate's COM and its constituent particles. Velocity #2 is updated similarly: each particle in an aggregate whose COM crosses a radial boundary receives a uniform velocity #2 offset immediately.

As noted in Quinn et al. (2010), the canonical momentum, P_y, must also be updated whenever a boundary crossing occurs, as the particles' angular momentum changes during this procedure. In our implementation, we find it is only necessary to store one P_y, corresponding to position and velocity pair #1.

However, once we introduce non-redundant velocities and positions for each particle, we immediately encounter a problem: how do we construct the other types of position and velocity from any given data? This issue occurs whenever particles are added to, or removed from, an aggregate—since the COM position immediately becomes unknown—and when initializing from stored data.

Our strategy begins with position and velocity pair #1, since that is the data stored in `pkdgrav` data files. Initially, the position and velocity pair #2 for all particles is set equal to #1. Then for each aggregate, we find the most massive particle, and use its position as a reference point (since that particle is most likely to be closest to the COM). Next, the code examines each particle in the aggregate, and determines if that particle's position #2 is currently more than half a patch length or width from the reference point. If it is, then it must be a wrapped particle (i.e., the yellow particle in the circle in Fig. 2), so its position #2 is adjusted (while updating velocity #2) until the particle is closer than half a patch dimension from the reference point (i.e., placing its position #2 within the circle).

Once all particles in the aggregate are unwrapped in this way, we calculate the aggregate's COM position based on the newly unwrapped positions #2. If the resulting COM position lies outside of the patch (which is a distinct possibility, given the arbitrary reference point), then the COM is wrapped by adjusting each constituent particle's position and velocity #2 accordingly, until the COM lies inside the patch.

We made a critical assumption in order to make this data reconstruction method possible: we assumed that aggregates will never grow to be larger than half a patch dimension across. Otherwise, the procedure will begin to mistake a distance from an aggregate's COM as large enough to require a wrap, when in fact this distance is real. Early testing shows that obvious numerical artifacts appear when aggregates grow large enough to break this procedure, and thus it is easy to discover when this problem manifests. If such a situation occurs, we must simply make the patch larger.

2.3.3. Aggregate orientation

Here we discuss how placing rigidly bonded aggregates into a non-inertial frame affects aggregate orientation.

The Eulerian equations of solid-body motion (Eqs. (3) and (4)) take an aggregate's spin and torque and update the body's spin rate and orientation of its principal axes. But these equations require the aggregate's principal axes, spin, and torque to all be in the body frame. We follow the rules in Section 2.3.1 to place each vector into the body frame for computation, and back into the patch frame for usage.

In addition, since our patch coordinate system rotates, sliding out of sync with the inertial space frame, we must carefully manipulate each aggregate's orientation. In an inertial frame, free of torques, an orbiting aggregate will maintain the orientation of its angular momentum vector. So, in the rotating frame, our aggregates must precess at the orbital frequency of the patch (Ω).

We handle this orientation issue when transforming to and from body coordinates before applying Eqs. (3) and (4). Consider an integration over the time interval dt. First rotate the aggregate into the space frame by rotating the transformation matrix $\mathcal{A}$ by an angle $+\Omega t$ (where t is the time since the space and patch frames were last in sync), and then go into the body frame by the rules in Section 2.3.1. When the Euler equations are done, transform from the body frame to the patch, and derotate $\mathcal{A}$ into the patch frame's new orientation, using: $-\Omega(dt + t)$.

In practice, we simplify this by arguing that the rotation angle $+\Omega t$ is arbitrary; so let $t = 0$ for all times in this exercise. Thus the angle by which we rotate into the "space" frame is zero degrees. This saves computational work. Thus the above procedure reduces to simply rotating $\mathcal{A}$ by $-\Omega dt$ after the Euler equations are finished.

This procedure results in an aggregate precessing clockwise (in the absence of external forces, such as tides), as viewed from the $\hat{z}$ direction in the patch. This is the proper behavior.

It is appropriate to note here that we have chosen to store the particle spins in the space frame, while the velocities of those particles (including velocity due to potential aggregate spin) are valid in the patch frame. This choice is made so that the particles will properly precess, as discussed above, while minimizing computation. If we stored the spins in the patch frame, we would have to rotate the spin vectors of perhaps millions of particles by $-\Omega dt\hat{z}$ at each timestep. This is more computational work than adjusting the spins as needed (i.e., during collision handing, after a collision is detected).

We recognize the inconsistency between our treatment of particle and aggregate orientation. This is done because aggregates are highly irregular in shape, and the proper aggregate orientations are needed in the patch frame multiple times during every step in order to properly compute the gravity field around an aggregate, and detect collisions. On the other hand, particles are perfect spheres, and their orientations do not affect gravity or collision detection. There will always be (far) fewer aggregates to update in this manner than particles, so this compromise avoids significant computational overhead.

2.3.4. Aggregate acceleration, torque, and stress

Here we discuss required modifications to aggregate acceleration, torque, and stress calculations when in a non-inertial frame. We compute these quantities for an aggregate using the accelerations on each of its constituent particles, applying the results of the derivations presented in Appendix A, as follows:

The acceleration of the COM of an aggregate is found by first computing a weighted sum of the accelerations due to gravity and external potentials of its constituent particles, and then computing the Hill equations of motion (Eq. (1)) for the COM. Thus the COM is accelerated by interparticle gravity, external potentials, and the Hill equations, acting dynamically as though it were a single rigid particle (see Eq. (30)).

In contrast, to compute both torques and stresses on rigid aggregates in the rotating frame, we require the accelerations on each particle without Hill's terms—since those add accelerations that do not torque a body directly. We achieve this by computing the torques and stresses after interparticle gravity and external forces are found, but before computing the Hill terms.

When the proper particle accelerations are known, we use Eq. (36) to calculate the torque on an aggregate. Note that we take care that all positions and the aggregate's inertia tensor are in the patch frame before applying this equation. The resulting torque vector then needs to be transformed into the body frame for use in the Euler equations, as mentioned above. Computed in this way, the planetary tides torque non-spherical aggregates to align with the radial axis of the planet.

Computing the stress on an aggregate determines if any of its rigid bonds should break (via stress fragmentation). Such fragmentation triggers include interparticle gravity, planetary tides, and rapid aggregate rotation. Also, if bonding is sufficiently weak, an irregularly shaped aggregate will fragment itself through self-gravity, and drive toward an equipotential shape (e.g., a sphere, if non-rotating).

We compute the stress on each particle in an aggregate separately, computing the relative acceleration between that particle and its aggregate's COM, including planetary tides (cf. Eq. (22)), and the centrifugal force from aggregate spin. As in Richardson et al. (2009), when that relative acceleration exceeds the tensile or shear strength (as defined by the user), the particle is liberated from the aggregate.

2.3.5. Collisions

Collision handling in the rotating frame, for the most part, remains unchanged from the methods used in Richardson (1995) and Richardson et al. (2009). But we must take care to properly include the reorientation of the aggregates due to the changes in the orientation of the rotating frame (see Section 2.3.3).

We have included the rotation of the patch frame in the velocity used for collision prediction. As in Richardson et al. (2009), collision detection for aggregates is performed by using a second-order prediction of each constituent particle's path (referred to as q in Richardson et al. (2009); see paragraph preceding their Eq. (A.4)), which includes considerations for aggregate spin. We add patch rotation to aggregate spin to obtain an estimate for each particle's trajectory. This provides an inexact approximation to the particle's true curved path, but it is accurate enough in short time intervals.

Once a collision is detected, and the collision time is found, we integrate the aggregate(s) involved forward to the collision time. This involves rotating $\mathcal{A}$ by the angle $-\Omega dt$ to compensate for frame rotation, integrating the Euler equations for aggregate spin, and drifting the COM. We then resolve the collision, including any potential merging or fragmentation events. Ideally, we would then trace back the COM position and derotate $\mathcal{A}$ into the proper orientation for the start of the step, in order to search for subsequent collisions. Unfortunately, that is not possible. The Euler equations that integrate the aggregate's orientation are computed with a time-adaptive Runge–Kutta integrator, which is not time-reversible.

Thus, the aggregate's COM is never "back drifted"—that is, after the collision is resolved, the aggregate's data is not returned back to a state that is valid for the beginning of the step. In order to keep track of when the aggregate's data is valid, each aggregate stores the last time it was updated—that is, the time at which its current position and orientation is valid. However, we must check for further collisions, and `pkdgrav`'s collision-search algorithm assumes that all particles positions are valid for the start of the step. Thus the positions for the aggregate's constituent particles are back drifted (undoing the frame rotation as well)—but the aggregate's COM is not.

Once all collisions have been resolved, aggregates finish the step by integrating forward over whatever time remains in the step (taking an Euler step, drifting the COM position, and rotating $\mathcal{A}$). If there are no collisions for an aggregate during a timestep, it simply takes a full step forward to complete the timestep.

2.4. Overlapping particles, and resolution strategies

As discussed in Richardson et al. (2009) and Section 2.3.5, collision detection for aggregates is done approximately. To produce a

predicted collision time, we extrapolate the rotation and motion of aggregates assuming, for example, that the spin rate of the aggregate does not change over the time interval being considered. However, integration forward to the predicted moment of impact is performed using Eqs. (3) and (4) with a Runge–Kutta integrator. Once the aggregates' orientation and position are accurately determined for the predicted collision time, the colliders may in fact not touch, or may find that they are overlapping. This is an unfortunate inevitability in this strategy—even a quartic expansion of the collision–prediction equations cannot provide the exact collision circumstances (though smaller timesteps help to reduce errors).[4]

We provide three options to account for particle overlaps, which we call the backstep, adjust position, and repel methods. Backstep calculates how far in the past the collision should have happened, moves the particles back (using their current velocities) to the point of impact, resolves the collision, and finally moves the particles to their new positions for the current time. This strategy works very well in simulations without aggregates; however, due to the irreversible nature of a Runge–Kutta integrator, the aggregates cannot be integrated back in time accurately in order to process the missed collision. Thus using this method with aggregates tends to result in numerical artifacts, such as abrupt changes in aggregate orientation.

The adjust position method simply moves the two overlapping particles apart (along their lines of centers) until they are just touching. This deceptively simple strategy also does not work with aggregates, for two reasons. The first is that the method does not move aggregates—it only moves the particles. For example, fixing an overlap in this way on the surface of an aggregate may cause a particle to move deeper into the aggregate, causing a host of additional unphysical overlaps within the aggregate. The second is that the adjust position strategy does not change the velocity of the particles, only the positions—thus angular momentum of aggregates is not conserved. Simulations have shown that this method causes aggregates to spin up spontaneously.

Finally, the repel method applies a repulsive force to overlapping particles such that they separate in a reasonable time. This force replaces the usual self-gravity between these overlapping particles (but other particles' gravity is calculated as usual). This repulsive force grows linearly with the degree of overlap, so particles barely overlapping feel a gentle push, while greatly overlapping particles feel a stronger force. The user controls how strong this force is via a parameter, called the repel factor. We note that higher repel factors mimic bouncier coefficient of restitution laws, as overlapping particles tend to separate at higher speeds. This repel method differs from the other two strategies in that it does not attempt to resolve the overlap instantaneously; rather it applies a gentle force that encourages particles to separate in time.

We have performed a suite of tests using the repel strategy (the only strategy we have that avoids numerical artifacts): Seventy to ninety percent of all overlaps are between free particles, and are highly transient. On average, less than 1% of the volume of any aggregate will be overlapped. Outliers do occur however, with some aggregates containing 10–20% of their volume in an overlap. However, these aggregates tend to be "dumbbells"—aggregates with only two particles—and these are easier to force apart than larger aggregates with more inertia. Also of note: roughly half of all overlaps between particles in aggregates are in fact in the same aggregate. This is a direct consequence of the inaccuracy of collision–prediction. These overlaps are frozen into the aggregate, since the particles cannot move relative to one another, and thus are not correctable (via the repel method) until some external process fragments the bond. We do not feel these overlaps are harmful to the results of our model, as the overlap simply imposes a larger-than-average mass density to that region of the aggregate, which mimics a physical variation in material density.

[4] An iterative procedure to predict and revise collision times may reduce these errors—a possible future feature of the code.

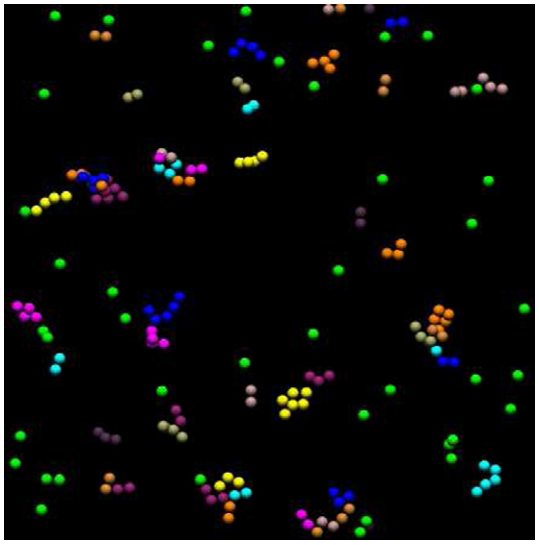

Fig. 3. Snapshot from a test simulation containing 200 particles in the coherent particle sliding-patch model, with merging and fragmentation. The particles are all 1 m in radius. Green bodies are free (unbonded) particles; when two free particles merge, they become an aggregate, and are assigned a random (non-green) color. When two aggregates merge, the particles in those aggregates do not change color. Many aggregates are visible in this frame; some are merely bonded pairs ("dumbbells"), while others are collections of many colors—indicating a complex aggregate–aggregate merging history. Viewing geometry is as in Fig. 1. Note that as in all simulations in this paper, this is a 3D model, so the appearance of major particle overlap is a projection effect.

3. Tests and early results

3.1. Small-scale testing

To verify that the model functions correctly, we ran small tests with 2–5 particles, confirming, for example: gyroscopic behavior as observed in the patch frame (precession at Ω); tidal forcing from the planet torquing aggregates to line up with the radial axis; spin-induced fragmentation (demonstrating that liberated fragments immediately move according to the Hill equations); merging particles conserving angular momentum in the non-inertial patch frame; and collision detection and resolution function in the rotating frame.

We note that our stress fragmentation model, as implemented, liberates particles from a spinning aggregate from the surface first, rather than from the center. (In a homogeneous solid, the greatest stress is at the center, due to the accumulated tensions of the surrounding accelerated material.) As noted in Section 2.1, our model does not maintain a bonding network, and thus cannot model transmission of tension across a solid. We recognize this limitation of the model, and further improvements to our model and technique (both discussed in Section 4) should mitigate this issue.

Fig. 3 shows a snapshot from a 200 particle simulation using the new cohesive particle model, illustrating the fractal nature of these small aggregates.

3.2. Large-scale testing

We next performed a suite of seventeen full-scale simulations with a range of merge limits and strengths. Each had 75,000

particles, and used nominal parameters for Saturn's outer A ring—namely, we placed the patch center 136,530 km from Saturn (with an orbital period of 14.3 h), and used a rather high estimated mass surface density of 500 kg/m^2 (in order to favor a higher collision rate), corresponding to a dynamical optical depth of 0.75. Our monodisperse particles were 1 m in radius, with an internal density of 0.5 g/cm^3 (to represent porous ice). The most significant departure from currently estimated A ring values is that our smallest dynamical unit was an indestructible sphere 1 m in radius. The model is capable of modeling a size distribution of indestructible particles, but in order to keep the number of particles tractable for these tests, we used 1 m monodisperse particles. (Note that a size distribution of bodies larger than 1 m emerges as the particles coagulate into aggregates.)

The dimensions of the patch were set to 4 by 10 λ_{crit} (approximately 880 by 350 m)—where λ_{crit} is the critical wavelength of the ring material (see Section 1.2)—with the longer axis in the azimuthal direction. The size of the patch must be large enough to resolve the largest structures in the patch, which on these scales are the gravity wakes (cf. Section 1.2), and also limit the effect of self-interaction between structures within the patch and their counterparts in the external replicated patches. The aspect ratio of 2:5 is chosen to accommodate the pitch angle of the gravity wakes. Following Porco et al. (2008), we used the velocity-dependent coefficient of restitution (ϵ) law of Borderies et al. (1984), using a v^* of 0.001 cm/s. This law has ϵ decrease with increasing impact speed, and our choice of v^* makes the law extremely dissipative; all impacts with speeds above 2 mm/s use an ϵ below 0.1, and the typical impact speed of 0.5 mm/s (the escape speed from a particle with the above parameters) has an ϵ of approximately 0.2.

Particles were placed randomly within the patch, with no initial aggregates, in a uniform vertical distribution 20 m thick. Initial velocities were chosen from a uniform distribution in each Cartesian coordinate, with minimum and maximum values of $\pm 2\Omega r$ in the radial and azimuthal directions, and $\pm\Omega r$ in the vertical direction (with $r = 1$ m). (In this way, we deliberately begin out of equilibrium, but note that—regardless of these choices—both the vertical scale height and the velocity dispersion of the particles equilibrate to consistent values within a few orbits.) Collisional merging and fragmentation were enabled, as was fragmentation due to stress. Each run was evolved for 15 orbits, which is sufficient to reach equilibration of the velocity dispersion. At this distance from Saturn, the orbital period is 14.3 h. Our simulations completed in 8–11 days on 16 processors, depending on the interparticle collision rate.

The critical opening angle (θ_{crit}) for the gravity tree was 0.5 radian, and our timestep was approximately 5 s. We tested a range of θ_{crit} values from 0.25 to 1 radian, and found no discernible systematic difference in outcome (besides a nearly a factor of two difference in computation time), and chose 0.5 as a conservative value. In testing a range of timesteps from approximately 0.5–50 s, we find our timestep is a reasonable compromise between accuracy and computation time. (A timestep of 50 s resulted in significant errors in the model's outcome; on the other hand, a timestep of 0.5 s yielded fewer overlaps, but otherwise no systematic differences, at the cost of over 10 times greater computation time.)

We studied two bonding parameters in two suites of nine simulations (with one simulation shared between each suite). The first suite studied the merge limit over the range of 0.01–1.0 v_{esc} (the escape speed of an individual particle, or 0.53 mm/s). The fragmentation limit was held constant at 1.0 v_{esc}. Thus the merge and fragmentation limits are constant among all collisions throughout each simulation, as we have a monodisperse population of particles. This choice is consistent with Hatzes et al. (1991), who found that below a constant impact speed of ~0.3 mm/s, ice spheres always tended to stick on impact.[5] This speed is 0.57 times the escape speed of our particles, which places it in the center of our parameter space. For this suite, we chose a constant, size-independent strength parameter of 100 Pa (for $\beta = 2$; see Appendix B), which is very roughly consistent with our choice of fragmentation limit (see Appendix C). This strength places this suite securely in the collision-dominated regime for fragmentation, as opposed to fragmentation via tides or rotational stresses (this assertion is demonstrated by the results below; see Fig. 8).

Note that, in the context of `pkdgrav` parameters (i.e., $\beta = 2$), a strength of 100 Pa approximately corresponds to 10^5 Pa in the frost layer overlap cohesion model (in which $\beta = 10^{-3}$). This distinction is necessary because of the differences in the contact areas between cohesive particles in different models (see Appendix B). For these tests, we make the assumption that strength is independent of the size of the body.

Our second suite of nine simulations studied the strength limit (in both tensile and shear) logarithmically over a parameter range of 10^{-5}–10^2 Pa (for $\beta = 2$), with one comparison run with infinite strength.[6] This suite held the merge limit fixed at 0.5 v_{esc}, or 0.27 mm/s, which closely matches the Hatzes et al. (1991) critical sticking velocity. All other parameters (fragmentation limit, patch size, etc.) were identical to the first suite.

Fig. 4 shows a snapshot of a simulated ring patch from this suite (see caption for parameter details). Gravity wakes are plainly visible in this image, mixed in with the bonded aggregates. Aggregates form primarily in the dense gravity wakes; this is expected, as the particle density is high, and relative speeds are low. The wakes shear out and disperse on timescales of less than an orbit, usually releasing a line of small non-spherical aggregates. These aggregates can then collide, both with dispersed wake material and other aggregates, and either merge and grow or fragment and erode, depending on the circumstances. Surviving aggregates eventually collide with other dense gravity wakes, at which time (given our choice of fragmentation limit) they are typically destroyed down to their constituent particles. That material is mixed into the wakes, which form new aggregates. However, in rare cases, larger aggregates colliding with wakes may erode somewhat before forming the core of a new large aggregate. This usually only occurs when the merge or fragmentation limits are set to very high levels (above v_{esc}). The vast majority of aggregates have a typical lifetime of somewhat less than an orbital time.

A particular quantity of interest in these simulations is the equilibrium aggregate size distribution. We estimate the size (effective radius r) of an irregularly shaped aggregate by first finding its principal axis lengths ($a \geqslant b \geqslant c$) then computing the radius of the volume-equivalent sphere as $r = \sqrt[3]{abc}$. Fig. 5 shows the incremental aggregate size distribution at the conclusion of one of the runs in the suite (using a merge limit of 0.5 v_{esc}). The distribution fits well to a power-law, $n(a) = n(a_0)[a_0/a]^{\alpha}$, with $\alpha = 3.0 \pm 0.23$. (The error quoted is the formal 1-σ uncertainty for a linear least-squares fit to the power-law.)

Fig. 6 shows the temporal variation of aggregate properties during a simulation, after the particles have settled into a reasonable equilibrium. There are no obvious long-term trends in these properties—indicating that the population has reached equilibrium—however significant variation is evident, as aggregates go through periods of heavy formation and subsequent destruction (usually when large gravity wakes collide). This plot demonstrates that

[5] We recognize that we are using particles forty times larger than those in Hatzes et al. (1991), and that size may have an effect on sticking properties, but this is not an important issue for these tests.

[6] Note that the first suite has no corresponding simulation with infinite merge limit; with our chosen parameters, such a setting is known to cause total collapse of the particles into a single aggregate.

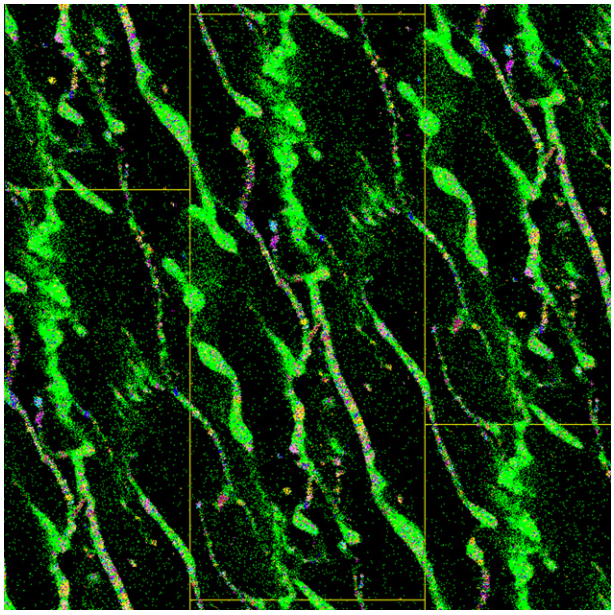

Fig. 4. Snapshot from a 75,000 monodisperse particle simulation with parameters consistent with Saturn's outer A ring. The patch (yellow box) is approximately 880 m long (azimuthal) and 350 m wide (radial). Aggregation parameters have the merge limit at 0.5 v_{esc} (particle escape speed, v_{esc}, is 0.53 mm/s), the fragmentation limit at 1.0 v_{esc}, and a bond strength (tensile and shear) of 100 Pa (for $\beta = 2$; see Appendix B). (See text for all remaining parameters.) Note that this merge limit approximately matches the critical sticking speed of 0.3 mm/s found by Hatzes et al. (1991). Colors and viewing geometry are as in Fig. 3. To enhance aggregate visibility, aggregates are drawn over non-bonded particles (green particles)—thus, aggregates are never hidden behind a cluster of green particles. Gravity wakes are clearly visible, and strings of aggregates appear throughout the image, at approximately the same pitch angle as the wakes; this is expected, as these aggregates are the remnants of previously dissolved wakes. (A particularly large wake, that is collapsing into a string of aggregates, is below-center in the image.) Many of these aggregates survive to collide with other wakes—examples of those aggregates can be seen freely floating in the interwake regions. However, the vast majority of aggregates are destroyed in less than an orbital time (<14.3 h) when they encounter another wake.

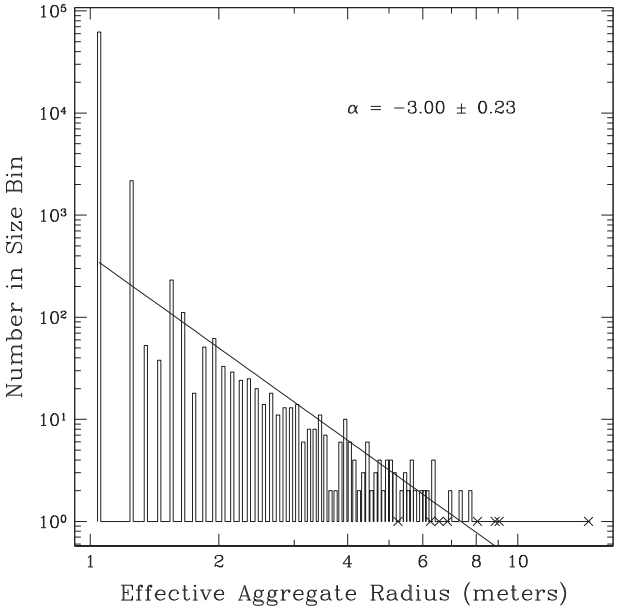

Fig. 5. Incremental size distribution of effective radii of bodies formed in an equilibrated region of the A ring in our cohesive aggregate patch model. This data is selected from the same simulation presented in Fig. 4 (see that figure caption for simulation parameters). Bins marked with an "x" have exactly one aggregate in them (since logarithmic axes cannot natively distinguish between bins containing zero or exactly one object). Assuming the function $n(a) = n(a_0)[a_0/a]^{\alpha}$, a linear (least-squares) fit to this log–log histogram (with bins 0.1 m wide) gives a power-law exponent $\alpha = -3.0 \pm 0.23$, shown with 1-σ uncertainty. The vast majority of bodies on the plot lie in the $r = 1$ m bin; these are the unbonded particles. The plot extends to include the largest aggregate at this moment in the simulation, with a radius of 14.7 m.

time-averaging helps to filter out these rapid changes in aggregate evolution.

For each simulation, we take the data for the final nine orbits and time-average them, producing the data displayed in Figs. 7 and 8. The time-variability of each property is indicated by the size of each errorbar; this is simply the standard deviation of each property's value over time. (Note that the error for α is computed slightly differently; see Fig. 6's caption.)

Fig. 7 shows the results from our first test suite, demonstrating that the aggregate populations that emerge vary significantly with the choice of merge limit. With higher merge limits (stickier particles) fewer particles remain unaggregated, and more aggregates form. In fact, the upper left pane, showing the percentage of unaggregated particles, is nearly a linear trend with merge limit, and extrapolates to nearly 100% at a merge limit of zero, as expected.

The right panels follow more intriguing trends. While increasing the merge limit increases both the α and the maximum aggregate size monotonically, much of our chosen range of merge limit ($\sim$0.2–1 v_{esc}, or 0.1–0.5 mm/s) produces an α of approximately -3, and a maximum effective radius of 10–20 m (within error bars). This is significant as these values are consistent with measurements of Saturn's A ring (Marouf et al., 1983; Zebker et al., 1985).

Fig. 8 shows the results of our second suite, in which we vary the strength (both tensile and shear) of the aggregates. (Note strengths in this figure are within the `pkdgrav` context, for $\beta = 2$; see Appendix B.) The leftmost side of each pane, where aggregates are weakest, shows the expected trends: fewer aggregated particles, fewer and smaller aggregates, and steeper size distributions. However, at a critical strength (S_{crit}) of $\sim 10^{-2}$ Pa (for $\beta = 2$), this trend ends. Above this critical strength, the results are nearly constant (any variation is owing to the stochastic nature of our simulations). In fact, the results of strengths from S_{crit} to infinity are roughly indistinguishable, indicating that throughout this strength range, fragmentation from stress is not occurring at a significant enough level to influence the emerging aggregate population. Thus, this result defines which fragmentation mechanism is dominant in which strength range: below S_{crit}, fragmentation is dominated by stress breakage (i.e., acceleration exceeding the user-defined strength); above S_{crit}, fragmentation is dominated by collisions (i.e., particle impact speeds exceeding the user-defined fragmentation limit). In the overlapping frost layer cohesion model, $S_{crit} \sim 10$ Pa (for $\beta = 10^{-3}$).

Our first suite used a fixed strength of 100 Pa (for $\beta = 2$), which is above S_{crit}. This shows that this suite was performed in the collision-dominated fragmentation regime, and that fragmentation via tides or aggregate rotation was not a factor in the results of the suite. Also of note: infinite strength cannot result in total collapse of the particles into a single aggregate, but extremely weak strength can prevent aggregation; strength = 10^{-5} Pa (for $\beta = 2$) resulted in a maximum average aggregate size of 2 m (two attached 1 m particles), but infinite strength only $\sim$20 m.

These two suites provide an excellent demonstration of the utility of this model. By simply varying two parameters, we have already discovered important behaviors of the aggregate population within our parameter space. These results will be expanded in future work as we explore more parameters beyond those demonstrated here.

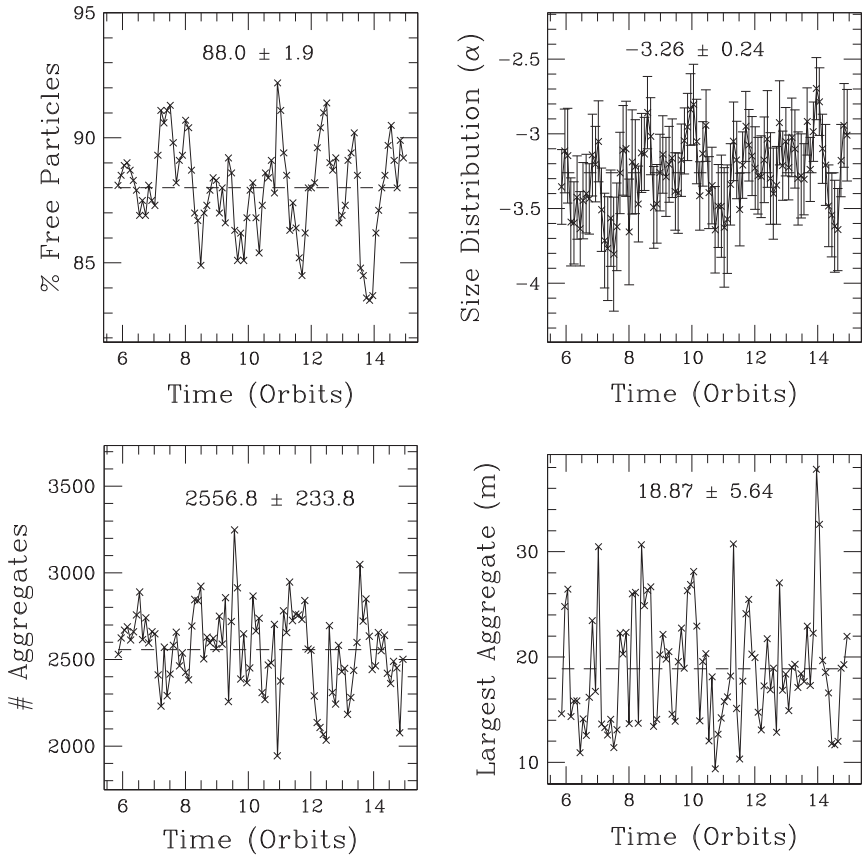

Fig. 6. Here we show the time evolution of four properties of the aggregate population in the same simulation as Fig. 5. This plot displays the final 9 orbits of the 15-orbit simulation, demonstrating the system's highly variable equilibrium. The upper-left panel is the fraction of unbonded particles in the simulation, with a value of 100% indicating no bonded particles at that time. The lower-left panel is the total number of aggregates in the simulation (the scale of which is of course dependent on the size of the simulated region). The upper-right panel shows the slope of the size distribution, α, taken from plots such as Fig. 5. The lower-right panel shows the effective radius of the largest aggregate, in meters. Note that the identity of the largest aggregate is not consistent throughout the simulation; in fact, the largest aggregate will change quite often, as aggregates are continually created and quickly destroyed, often existing for less than one orbit. The time-averaged data is displayed on each pane as a mean and a 1-σ standard deviation, with the mean drawn on the plot as a dashed line. The error for α is more complex than a simple standard deviation of its variation in time, as each point has its own error (from the least-squares fit). Propagation of error was used to combine each point's error with the variance in the time-averaged data to obtain the value displayed on the plot; however, we note that most of the quoted error in α results from time-variability.

4. Summary

We have combined two numerical methods—a local patch model (e.g., Porco et al., 2008), and a cohesive aggregate model (e.g., Richardson et al., 2009)—forming a new tool with unique capabilities. The final result is a local *N*-body code in a rotating coordinate system, with self-gravity and hard-sphere collisions, that includes self-consistent rigid-body cohesion and fragmentation. This model is useful in building a first-order understanding of any system that involves a large number of colliding and cohesive bodies, each far smaller than the size of the system, orbiting a central body—e.g., ring dynamics and planetary system formation. However, the cohesive aspects of our *N*-body model remain quite simple, and while potentially applicable to a wide range of bonding methods, our approach does not attempt to encapsulate the detailed granular mechanics (e.g., bonding networks) and fragmentation physics (e.g., production of fragments with realistic size distributions) of more sophisticated models.

We have detailed how we have modified the existing `pkdgrav` model's rigid-body capabilities to function in a local, rotating frame. We discussed the frequency and severity of particle overlaps in our model, and our resolution strategies, which we have found to be adequate for the physical requirements of the systems we are modeling.

Last, we provided a demonstration of this model at work, as we varied two of our modeled bonding parameters: the merge limit (velocity threshold for aggregation) and the strength (the aggregated particles' resistance to breakage via tidal, centrifugal, or gravitational stress). These example data suites showed that our bonding parameters—only one of which has an experimental constraint—have an important effect on the final state of the system. We demonstrated an example range of merge limits that results in observable properties similar to those measured in Saturn's A ring by the Voyager spacecraft: ∼0.1–0.5 mm/s, which brackets the merge limit for icy spheres found experimentally (∼0.3 mm/s) by Hatzes et al. (1991).

We also identified two strength regimes within the parameter space examined: at extremely low strengths, fragmentation is dominated by stress breakage, while at moderately high strengths, fragmentation is dominated by interparticle collisions. On the scale of ∼10 m aggregates, the critical strength (S_{crit}) is approximately 10^{-2} Pa (assuming a contact area of a full particle cross-section—i.e., $\beta = 2$; see Appendix B), or ∼10 Pa in the overlapping frost layer cohesion model ($\beta = 10^{-3}$). (Note that these strength-related results are subject to many assumptions, and are accurate to an order of magnitude at best.)

As mentioned above, our model provides a simple means of experimenting with rigid aggregation in a rotating frame, but its simplicity does lead to subtle behaviors that are contrary to the physical behaviors of rigid solids. The most significant of these is that stress fragmentation occurs from the outside-in, rather than from the inside-out (that is, the model liberates particles from

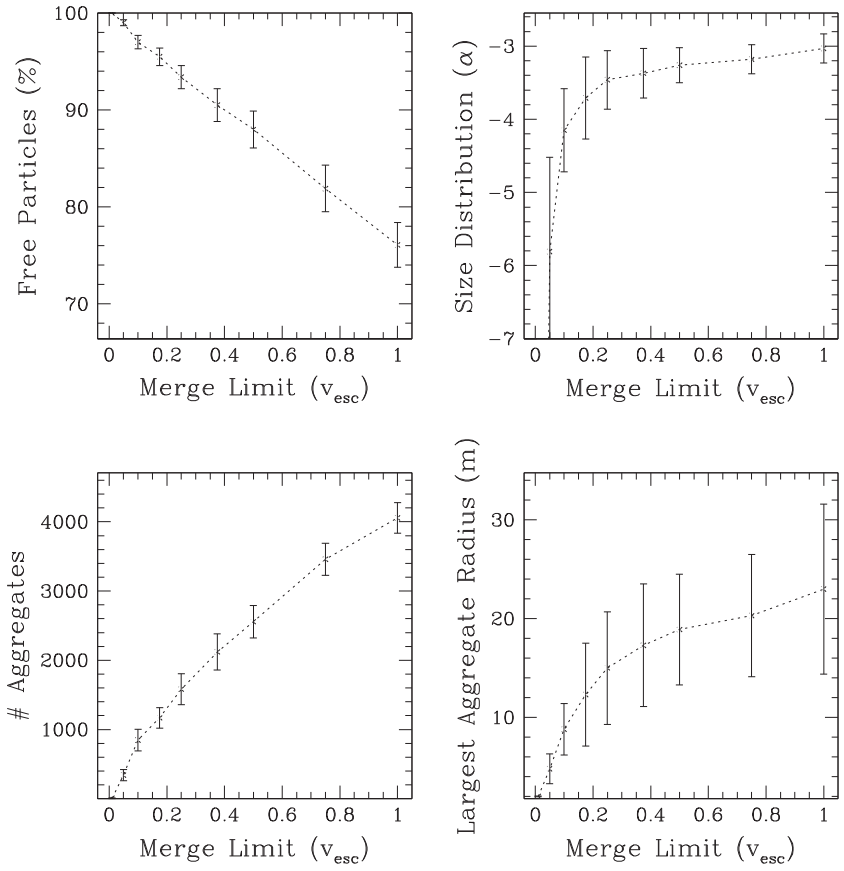

Fig. 7. Here we demonstrate the results of our first suite of nine simulations, in which the merge limit was varied. Displayed versus merge limit (defined in Section 2.1) are the same time-averaged properties shown in Fig. 6. Each point was obtained for each run as described in the caption for Fig. 6. As the merge limit rises, it is easier for colliding particles to merge, and this has a visible effect on the equilibrium aggregate population. Note the somewhat abrupt change in the behavior in α, and the radius of the largest aggregate, below a merge limit of approximately 0.2 v_{esc}—deviating from observed A ring properties from Voyager (Marouf et al., 1983; Zebker et al., 1985). For completeness: in the upper-right pane, the calculated value of α for a merge limit of 0.01 is −15.3 (far below the plot's limits), indicative of extremely limited aggregate formation.

the surface of an aggregate under stress, rather than breaking the aggregate into two large remnants). Were this behavior a dominant phenomenon in our simulations, our equilibrated models may favor smaller bodies, and under-represent large bodies, resulting in steeper equilibrium size distributions. However, we have demonstrated that our simulations can be carried out in a collisionally dominated fragmentation regime (so long as the strength parameter is above the critical strength, S_{crit})—thus preventing this phenomenon from significantly affecting our data. We recognize that the results of our model are less convincing near and below this critical value, and do not plan on studying such weak bonds with this model.

Further work will expand the testing in Section 3.2 from a demonstration to a full examination of the parameter space. The first aim will be to constrain the potential range of bonding parameters that continue to match the Voyager data of particles in Saturn's A ring, and then make predictions for the—yet unobserved—size distribution and potential maximum particle size in the B ring. This study will shed light on the hypothesis discussed in Tremaine (2003) regarding shear-free regions of the B ring, which estimated that a bond strength of approximately 10,000 Pa (10^5 dyn/cm^2) would be sufficient to allow particle assemblies of approximately 100 km to resist tidal shear. However, it is not known whether aggregates can grow to that size when dynamic collisional fragmentation is taken into account.

Further, it is possible to use this new tool to explore the theory proposed by Porco et al. (2008) that wake–interwake optical depth contrast is enhanced by significant damping. Cohesive bonding should have an overall effect on collisional damping that reduces the net coefficient of restitution—and can be quantified via optical depth variations using, for example, the analytical technique detailed in Tiscareno et al. (2010).

Several refinements to improve the realism of the cohesive model are under development. These refinements include attaching breakable springs between individual particle centers. As the particles separate, they feel a Hooke's law restoring force that provides a resistance to separation. Once a maximum distension is reached, the spring can break, freeing the attached particles. This model captures the desirable property of failure stress maximizing at the body center, but greatly increases the complexity of particle rotation tracking. Additionally, work to improve collision detection between aggregates may lessen the severity of particle overlaps in our model—however this appears to be a minor concern.

Acknowledgments

R.P.P. acknowledges support from the National Aeronautics and Space Administration under the NASA Earth and Space Science Fellowship program, and D.C.R. from Grant No. NNX08AM39G issued through the Office of Space Science. D.J.S. acknowledges support from NASA's Planetary Geology and Geophysics program.

The authors thank S. Tremaine for his thoughtful discussions on this topic, and two anonymous reviewers for their suggestions that improved the manuscript.

The simulations presented in this paper were carried out using computing clusters administered by the Center for Theory and

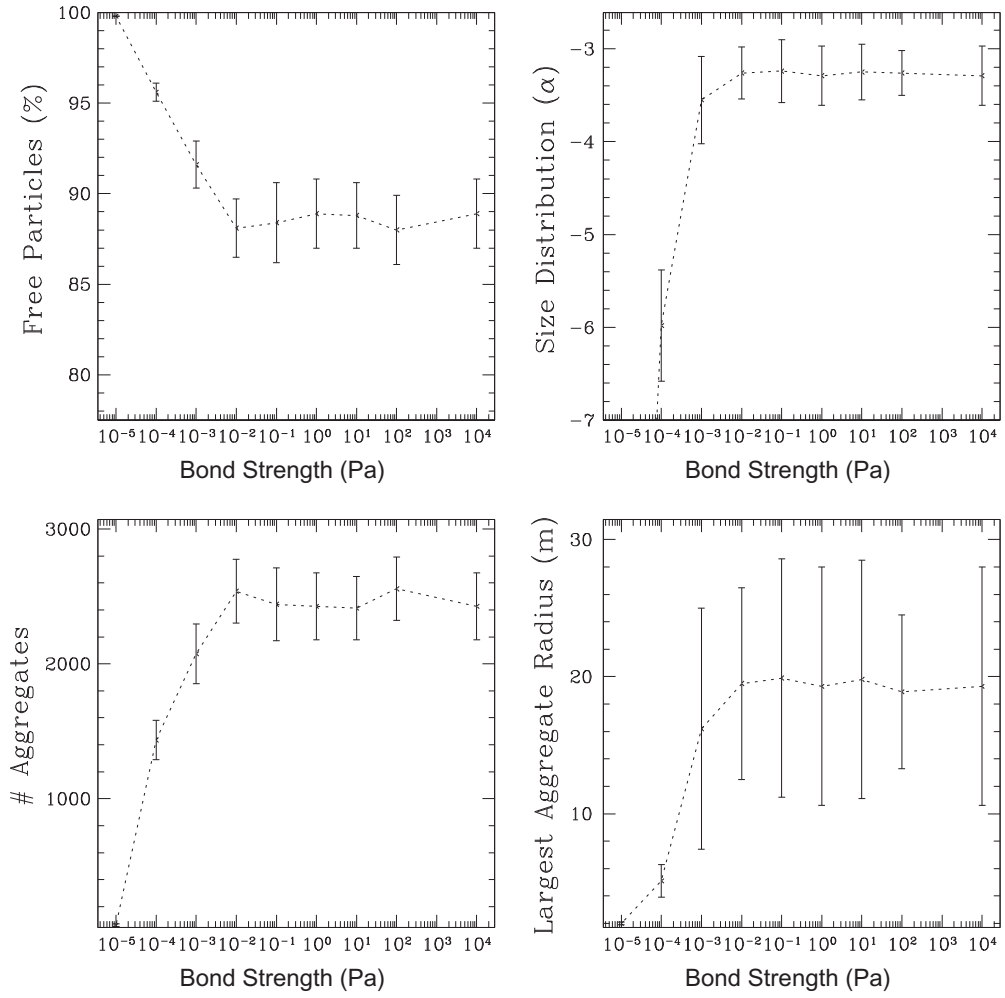

Fig. 8. Similar to Fig. 7, plotted here are the resulting changes in the aggregate population across our second suite of nine simulations in which the strength, or particle stress limit, was varied. Note that strengths are plotted on this figure in the context of the `pkdgrav` strength parameter (i.e., $\beta = 2$). As the strength rises, particles must experience larger accelerations relative to their aggregate's center of mass in order to separate from the aggregate. The simulation that is plotted at "10^4 Pa" had, in fact, infinite strength—that is, aggregates could never break due to stress. The results of all of the simulations—above an apparent critical value, S_{crit}—are indistinguishable from the case of infinite strength. This critical value is approximately 10^{-2} Pa (for $\beta = 2$). At weaker strengths, aggregates are smaller and less numerous, as aggregate spin and tides begin to cause aggregate breakage. At 10^{-5} Pa (for $\beta = 2$), aggregates can barely hold as 2-particle "dumbbells," which is reflected clearly in all panels of the plot. For completeness: $\alpha = -14.6$ at a strength of 10^{-5} Pa. As a reminder, strengths in `pkdgrav` are converted from separation forces in a very general way, using the particle cross-section as the full contact area (i.e., $\beta = 2$), and the reader should keep in mind the discussion in Section 2.1 and Appendix B when interpreting these results in the context of a particular cohesion model.

Computation of the Department of Astronomy at the University of Maryland ("YORP"), and the Office of Information Technology at the University of Maryland ("HPCC").

Appendix A. Derivation of the $N + 1$ Hill Problem with rigid bodies

A.1. Introduction

This derivation starts from the $N + 1$ body problem and derives the "$N + 1$ Hill Problem". Then, supposing that a subset of the masses are joined into a rigid body, it forms the appropriate gravitational force and torque acting on a rigid body in this problem. We start with the classical $N + 1$ body problem in an inertial frame. Note that we will remain in an inertial frame until the very end, as this makes the derivation simpler. Much confusion arises when a transformation into a rotating frame is made, as not all of the accelerations that arise can be disentangled easily from the forces acting on the bodies.

Consider $N + 1$ mutually gravitating bodies P_i with position vectors $\boldsymbol{R}_i$ measured from the system barycenter and masses m_i, $i = 0, 1, \ldots, N$. By Newton's laws their motion is defined by the equations of motion:

$$m_i \ddot{\boldsymbol{R}}_i = G \sum_{j=0, j\neq i}^{N} \frac{m_i m_j}{|\boldsymbol{R}_{ij}|^3} \boldsymbol{R}_{ij} \tag{5}$$

where $\boldsymbol{R}_{ij} = \boldsymbol{R}_j - \boldsymbol{R}_i$, and G is the gravitational constant.

We will distinguish the body P_0 from the rest of the system. To that end define $\boldsymbol{R}$ as:

$$\boldsymbol{R} = \frac{1}{m'} \sum_{i=1}^{N} m_i \boldsymbol{R}_i \tag{6}$$

where $m' = \sum_{i=1}^{N} m_i$. From the barycentric assumption we find that:

$$m_0 \boldsymbol{R}_0 + m' \boldsymbol{R} = 0 \tag{7}$$

or

$$\boldsymbol{R}_0 = -\mu \boldsymbol{R} \tag{8}$$

where $\mu = m'/m_0$.

We note that we can explicitly write out the equation of motion for $\boldsymbol{R}$:

$$\ddot{\boldsymbol{R}} = -G\frac{m_0}{m'}\sum_{j=1}^{N}\frac{m_j}{|\boldsymbol{R}_{0j}|^3}\boldsymbol{R}_{0j} \tag{9}$$

A.2. Shift of center

Now we will shift the center of the remaining N bodies relative to the vector $\boldsymbol{R}$, i.e., we will redefine the position vectors as:

$$\boldsymbol{R}_i = \boldsymbol{R} + \boldsymbol{r}_i \tag{10}$$

Under this change we note the following:

$$\boldsymbol{R}_i - \boldsymbol{R}_0 = \boldsymbol{r}_i + (1+\mu)\boldsymbol{R} \tag{11}$$

$$\sum_{i=1}^{N} m_i \boldsymbol{r}_i = 0 \tag{12}$$

Using the transformation $\boldsymbol{r}_i = \boldsymbol{R}_i - \boldsymbol{R}$ and the fact that we are still in an inertially oriented frame (i.e., a non-rotating frame), we can derive the equations of motion for the new positions as:

$$\ddot{\boldsymbol{r}}_i = \ddot{\boldsymbol{R}}_i - \ddot{\boldsymbol{R}} \tag{13}$$

$$= \frac{Gm_0}{m'}\sum_{j=1}^{N}\frac{m_j(\boldsymbol{r}_j + (1+\mu)\boldsymbol{R})}{|\boldsymbol{r}_j + (1+\mu)\boldsymbol{R}|^3} - Gm_0\frac{(\boldsymbol{r}_i + (1+\mu)\boldsymbol{R})}{|\boldsymbol{r}_i + (1+\mu)\boldsymbol{R}|^3} + \sum_{j=1\,j\neq i}^{N}\frac{Gm_j\boldsymbol{r}_{ij}}{|\boldsymbol{r}_{ij}|^3} \tag{14}$$

A.3. The Hill approximation

Now we will apply the Hill approximation, following Scheeres (1998). We note that the assumptions are simple, namely that $\mu \ll 1$ and $|\boldsymbol{r}_{jk}| \ll |\boldsymbol{R}|$. The heart of the approximation is the expansion of the denominator terms $|\boldsymbol{r}_i + (1+\mu)\boldsymbol{R}|^3$ and their multiplication with the numerator terms, ignoring higher powers of r/R (where $R = |\boldsymbol{R}| = \sqrt{\boldsymbol{R}\cdot\boldsymbol{R}}$). Specifically, we note the following:

$$\frac{(\boldsymbol{r}_i + (1+\mu)\boldsymbol{R})}{|\boldsymbol{r}_i + (1+\mu)\boldsymbol{R}|^3} = \frac{(\boldsymbol{r}_i + (1+\mu)\boldsymbol{R})}{(1+\mu)^3R^3\left[1 + \frac{2\boldsymbol{r}_i\cdot\boldsymbol{R}}{(1+\mu)R^2} + \left(\frac{r_i}{(1+\mu)R}\right)^2\right]^{3/2}} \tag{15}$$

which can be expanded using the binomial theorem to find:

$$\frac{1}{R^3}\left[\boldsymbol{R} + \boldsymbol{r}_i - \frac{3}{R^2}(\boldsymbol{R}\cdot\boldsymbol{r}_i)\boldsymbol{R} + \cdots\right] \tag{16}$$

where higher orders of μ have also been neglected. This is just a linearization of the force terms that depend on the distance R.

Substituting the above into our equations of motion yields:

$$\ddot{\boldsymbol{R}} = -\frac{Gm_0}{m'R^3}\sum_{j=1}^{N} m_j\left[\boldsymbol{R} + \boldsymbol{r}_j - 3(\hat{\boldsymbol{R}}\cdot\boldsymbol{r}_j)\hat{\boldsymbol{R}}\right] \tag{17}$$

$$= -\frac{Gm_0\boldsymbol{R}}{R^3} \tag{18}$$

where the other terms sum to zero given Eq. (12). Similar simplifications occur in the equations for $\boldsymbol{r}_i$, yielding:

$$\ddot{\boldsymbol{r}}_i = -\frac{Gm_0\boldsymbol{r}_i}{R^3} + \frac{3Gm_0}{R^3}(\hat{\boldsymbol{R}}\cdot\boldsymbol{r}_i)\hat{\boldsymbol{R}} + \sum_{j=1\,j\neq i}^{N}\frac{Gm_j\boldsymbol{r}_{ij}}{|\boldsymbol{r}_{ij}|^3} \tag{19}$$

Note that the equation for $\boldsymbol{R}$ is just the 2-body problem, and thus the center of mass of the N particles describes a Keplerian orbit. The simplest, and usual, case is to assume that it follows a circular path, as then the radius R is constant and the term $Gm_0/R^3 = \Omega^2$ and is the mean motion of the particle center of mass about the central body. Rewriting the equations in this way we find:

$$\ddot{\boldsymbol{r}}_i = -\Omega^2\boldsymbol{r}_i + 3\Omega^2(\hat{\boldsymbol{R}}\cdot\boldsymbol{r}_i)\hat{\boldsymbol{R}} + \sum_{j=1\,j\neq i}^{N}\frac{Gm_j\boldsymbol{r}_{ij}}{|\boldsymbol{r}_{ij}|^3} \tag{20}$$

Further, if we transform to a uniformly rotating reference frame, rotating with the mean motion and about the $\hat{\boldsymbol{z}}$ axis, we find the usual form of the Hill Equations, written in a rotating reference frame:

$$\ddot{\boldsymbol{r}}_i + 2\Omega\hat{\boldsymbol{z}}\times\dot{\boldsymbol{r}}_i = -\Omega^2\hat{\boldsymbol{z}}\times\hat{\boldsymbol{z}}\times\boldsymbol{r}_i - \Omega^2\boldsymbol{r}_i + 3\Omega^2(\hat{\boldsymbol{R}}\cdot\boldsymbol{r}_i)\hat{\boldsymbol{R}} + \sum_{j=1\,j\neq i}^{N}\frac{Gm_j\boldsymbol{r}_{ij}}{|\boldsymbol{r}_{ij}|^3} \tag{21}$$

$$= 3\Omega^2(\hat{\boldsymbol{R}}\cdot\boldsymbol{r}_i)\hat{\boldsymbol{R}} - \Omega^2(\hat{\boldsymbol{z}}\cdot\boldsymbol{r}_i)\hat{\boldsymbol{z}} + \sum_{j=1\,j\neq i}^{N}\frac{Gm_j\boldsymbol{r}_{ij}}{|\boldsymbol{r}_{ij}|^3} \tag{22}$$

where the $\hat{\boldsymbol{R}}$ direction is usually taken to be the x-axis. This is the appropriate set of equations to integrate to track the motion of the particles relative to their center of mass and incorporating mutual attractions between them. As such, Eq. (22) is used to calculate stress within an aggregate (see Section 2.3.4). Note that this is the generalized form of Eq. (1), providing the derivation for the Hill equations of motion both for single particles and aggregates in the N-body frame.

A.4. Forces and torques acting on a rigid body

Now assume that a subset of the N bodies P_i are rigidly tied to each other. Let us assume that the first L bodies, $P_1, P_2, \ldots, P_L$, with $L < N$, are tied together in a rigid body. Define the total mass of this body to be $m'' = \sum_{i=1}^{L} m_i$ and its center of mass to be $\boldsymbol{r} = \frac{1}{m''}\sum_{i=1}^{L} m_i\boldsymbol{r}_i$.

We then wish to find the total force acting on this body and the net torque acting on it about its center of mass. We should note that the total force that acts on a body P_i is computed using Eq. (20) as we are only interested in the gravitational forces acting on the particles. Once the forces and torques acting on the rigid body in the inertial frame have been stated, then we can transform to a rotating frame if desired.

A.4.1. Rigid-body center of mass equation of motion

Thus, the force acting on the body P_i due to mutual attraction of all the other bodies, under the Hill approximation, can be stated as:

$$\boldsymbol{F}_i = m_i\ddot{\boldsymbol{r}}_i \tag{23}$$

$$= -\Omega^2 m_i\boldsymbol{r}_i + 3\Omega^2 m_i(\hat{\boldsymbol{R}}\cdot\boldsymbol{r}_i)\hat{\boldsymbol{R}} + \sum_{j=1\,j\neq i}^{N}\frac{Gm_im_j\boldsymbol{r}_{ij}}{|\boldsymbol{r}_{ij}|^3} \tag{24}$$

The total force acting on the rigid body is then:

$$\boldsymbol{F} = \sum_{i=1}^{L}\boldsymbol{F}_i \tag{25}$$

$$= -\Omega^2 m''\boldsymbol{r} + 3\Omega^2 m''(\hat{\boldsymbol{R}}\cdot\boldsymbol{r})\hat{\boldsymbol{R}} + \sum_{i=1}^{L}\sum_{j=1\,j\neq i}^{N}\frac{Gm_im_j\boldsymbol{r}_{ij}}{|\boldsymbol{r}_{ij}|^3} \tag{26}$$

$$= -\Omega^2 m''\boldsymbol{r} + 3\Omega^2 m''(\hat{\boldsymbol{R}}\cdot\boldsymbol{r})\hat{\boldsymbol{R}} + \sum_{i=1}^{L}\sum_{j=L+1\,j}^{N}\frac{Gm_im_j\boldsymbol{r}_{ij}}{|\boldsymbol{r}_{ij}|^3} \tag{27}$$

where we note in the last step that the forces between the rigid body particles identically vanish. The equation of motion for the rigid body center of mass is then

$$\ddot{\boldsymbol{r}} = \frac{1}{m''}\boldsymbol{F} \tag{28}$$

$$= -\Omega^2\boldsymbol{r} + 3\Omega^2(\hat{\boldsymbol{R}}\cdot\boldsymbol{r})\hat{\boldsymbol{R}} + \frac{1}{m''}\sum_{i=1}^{L}\sum_{j=L+1\,j}^{N}\frac{Gm_im_j\boldsymbol{r}_{ij}}{|\boldsymbol{r}_{ij}|^3} \tag{29}$$

Again, if we transform into a rotating coordinate frame the equation takes on a more standard form:

$$\ddot{\boldsymbol{r}} + 2\Omega\hat{\boldsymbol{z}} \times \dot{\boldsymbol{r}} = 3\Omega^2(\hat{\boldsymbol{R}} \cdot \boldsymbol{r})\hat{\boldsymbol{R}} - \Omega^2(\hat{\boldsymbol{z}} \cdot \boldsymbol{r})\hat{\boldsymbol{z}} + \frac{1}{m''}\sum_{i=1}^{L}\sum_{j=L+1,j}^{N}\frac{Gm_i m_j \boldsymbol{r}_{ij}}{|\boldsymbol{r}_{ij}|^3} \tag{30}$$

A.4.2. Rigid-body torques

Computation of the torques is a bit more involved, yet yields a very simple result. The torque from each particle P_i relative to the rigid-body center of mass $\boldsymbol{r}$ is computed as $\boldsymbol{N}_i = (\boldsymbol{r}_i - \boldsymbol{r}) \times \boldsymbol{F}_i$. It can be verified that the torque equals the following:

$$\begin{aligned}\boldsymbol{N}_i &= 3\Omega^2 m_i(\boldsymbol{r}_i - \boldsymbol{r}) \times \hat{\boldsymbol{R}}(\hat{\boldsymbol{R}} \cdot (\boldsymbol{r}_i - \boldsymbol{r})) - \Omega^2 m_i(\boldsymbol{r}_i - \boldsymbol{r}) \times \boldsymbol{r} \\ &\quad + 3\Omega^2 m_i(\boldsymbol{r}_i - \boldsymbol{r}) \times \hat{\boldsymbol{R}}(\hat{\boldsymbol{R}} \cdot \boldsymbol{r}) + m_i\sum_{j=1,j\neq i}^{N}\frac{Gm_j(\boldsymbol{r}_i - \boldsymbol{r}) \times \boldsymbol{r}_{ij}}{|\boldsymbol{r}_{ij}|^3}\end{aligned} \tag{31}$$

Summing this over the L bodies causes the two middle terms to vanish, leaving the result:

$$\boldsymbol{N} = \sum_{i=1}^{L}\boldsymbol{N}_i \tag{32}$$

$$= -3\Omega^2\hat{\boldsymbol{R}} \times \left[\sum_{i=1}^{L} m_i(\boldsymbol{r}_i - \boldsymbol{r})(\boldsymbol{r}_i - \boldsymbol{r})\right] \cdot \hat{\boldsymbol{R}} + G\sum_{i=1}^{L}\sum_{j=1,j\neq i}^{N}\frac{m_i m_j(\boldsymbol{r}_i - \boldsymbol{r}) \times (\boldsymbol{r}_j - \boldsymbol{r}_i)}{|\boldsymbol{r}_{ij}|^3} \tag{33}$$

where the term within the first summation is just the outer product of the two vectors.

Each of the terms in this equation can be simplified. For the first, we note that the inertia dyadic of any body with respect to its center of mass consisting of point masses can be stated as:

$$\boldsymbol{I} = \sum_{i=1}^{L} m_i[(\boldsymbol{r}_i - \boldsymbol{r}) \cdot (\boldsymbol{r}_i - \boldsymbol{r})\boldsymbol{U} - (\boldsymbol{r}_i - \boldsymbol{r})(\boldsymbol{r}_i - \boldsymbol{r})] \tag{34}$$

where $\boldsymbol{U}$ is the identity dyadic, $\boldsymbol{r}_i$ are the positions of the point masses, $\boldsymbol{r}$ is the center of mass of the body and $\boldsymbol{I}$ is the inertia dyadic. In Eq. (33) we note that the term involving the identity will vanish, meaning that we can substitute the full inertia dyadic in the above equation for the quantity within the summation.

For the mutual gravitational term we note the identity $\boldsymbol{r}_{ij} = \boldsymbol{r}_j - \boldsymbol{r} - (\boldsymbol{r}_i - \boldsymbol{r})$, allowing us to simplify the result and just focus on the terms that do not mutually cancel. Applying both simplifications we find the torque acting on the rigid body:

$$\boldsymbol{N} = 3\Omega^2\hat{\boldsymbol{R}} \times \boldsymbol{I} \cdot \hat{\boldsymbol{R}} + G\sum_{i=1}^{L}\sum_{j=1,j\neq i}^{N}\frac{m_i m_j(\boldsymbol{r}_i - \boldsymbol{r}) \times (\boldsymbol{r}_j - \boldsymbol{r})}{|\boldsymbol{r}_{ij}|^3} \tag{35}$$

Again the common terms from the rigid body will mutually cancel, leading to

$$\boldsymbol{N} = 3\Omega^2\hat{\boldsymbol{R}} \times \boldsymbol{I} \cdot \hat{\boldsymbol{R}} + G\sum_{i=1}^{L}\sum_{j=L+1}^{N}\frac{m_i m_j(\boldsymbol{r}_i - \boldsymbol{r}) \times (\boldsymbol{r}_j - \boldsymbol{r})}{|\boldsymbol{r}_{ij}|^3} \tag{36}$$

This is the torque due to the gravitational forces acting on the rigid body, accounting for the Hill approximation. This can be used in the Euler equations, so long as the torque is stated in the proper frame.

A.4.3. Rigid-body rotational equations

Given the body torques acting on the rigid body (Eq. (36)), and the translational equations that define its center-of-mass location (e.g., Eq. (1)), we can state the full rotational dynamics equations for this body. The fundamental equation that controls rotational motion is:

$$\dot{\boldsymbol{H}} = \boldsymbol{N} \tag{37}$$

$$\boldsymbol{H} = \boldsymbol{I} \cdot \boldsymbol{\omega} \tag{38}$$

where the time derivative is with respect to an inertially oriented frame, the inertia dyadic is evaluated at the body center of mass, $\boldsymbol{H}$ and $\boldsymbol{\omega}$ are the angular momentum and velocity vectors of the body relative to an inertially oriented frame, and $\boldsymbol{N}$ are the total torques acting on the rigid body relative to the center of mass. As the equation is a vector equation, we have freedom in expressing which frame they are relative to.

If these are expressed relative to a frame rotating with an angular velocity $\boldsymbol{\Omega}$ relative to an inertially fixed frame, the equations become more complex. First, we recall the result for the time derivative of a vector relative to a rotating frame:

$$\dot{\boldsymbol{x}} = \boldsymbol{x}' + \boldsymbol{\Omega} \times \boldsymbol{x} \tag{39}$$

where $'$ signifies the time derivative relative to the rotating frame. Applying this to the fundamental equation we find:

$$\boldsymbol{H}' + \boldsymbol{\Omega} \times \boldsymbol{H} = \boldsymbol{N} \tag{40}$$

This is the general derivation for the Euler equations (Eq. (3)), accounting for the Hill problem frame in the statement of the torque $\boldsymbol{N}$. Note that Eq. (3) is simplified for use in the body frame of the rotating aggregate.

Appendix B. Contact area for intersecting spheres

Here we derive an approximate expression for the contact area between two overlapping equal-size spheres. (This expression should additionally be applicable to deformable spheres, as noted.)

We refer to Fig. 9 and its caption to define the symbols R, R', d, and h. Physically, d represents the depth of particle overlap (or two times the deformation depth, depending on the model). Note that $d + R' = R$, and that h and $R' + 0.5d$ form a right triangle, with R as the hypotenuse. We assume that the circular contact area between these two equal-size spheres is πh^2. (This is an assumption in the case of overlapping particles, but should be satisfactory with deformable particles.)

By virtue of our right triangle, the Pythagorean theorem gives us:

$$(R' + 0.5d)^2 + h^2 = R^2 \tag{41}$$

Next, we remove R' by its relationship to d and R. Rearranging and canceling yields:

$$h^2 = Rd - d^2/4 \tag{42}$$

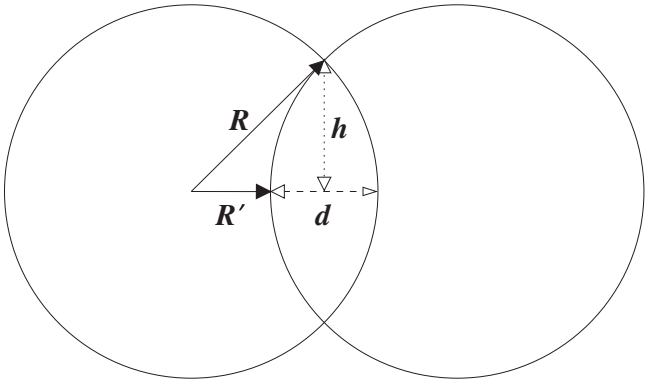

Fig. 9. Geometry and symbols used in the derivation of an approximate value for the contact area between intersecting spheres (Appendix B). Here, d refers to the dashed line, indicating the depth of particle overlap; h is the dotted line, indicating the distance from the center of the overlap to the edge of the particle intersection—in this approximation, h also refers to the radius of the circular particle contact area; R' is the shortest distance from a particle center to the intersecting particle; and R is the particle radius.

Now, we define a new parameter β:

$$\beta \equiv d/R \tag{43}$$

which can be understood as the fraction of a particle's radius that is overlapped (or deformed). This parameter encapsulates all of the unknown bonding physics into one general term.

In terms of this new parameter, the circular contact area is:

$$\pi h^2 = \pi R^2 \beta(1 - \beta/4) \tag{44}$$

When $\beta \ll 1$, this reduces to

$$\pi h^2 = \pi R^2 \beta \tag{45}$$

Thus, at small values of β, the contact area approximately equals the cross-section of the particles times β.

If $\beta = 2$, the contact area equals the cross-section, πR^2, and the two particles are completely overlapped, with their centers touching (or their surfaces are deformed such that their centers touch). And if $\beta = 0$, then there is no overlap or deformation, and the contact area is zero.

Contact area scales inversely with strength (with constant force). As such, Eq. (45) provides a simple means to convert from any strength used as a parameter in `pkdgrav` to a strength meaningful in whatever bonding model is under study. Since `pkdgrav` assumes that the contact area is πR^2 (consistent with $\beta = 2$ in Eq. (44)), it is a simple matter of finding the β value for the bonding model and dividing β into `pkdgrav`'s strength parameter (or $\beta(1 - \beta/4)$ if $\beta \sim 1$). For example, our bonding concept is based on overlapping frost layers; Hatzes et al. (1991) studied 2.5 cm ice spheres, and found that only "the outermost 10–50 μm of the frost layer [was] largely responsible for the sticking mechanism." Taking a convenient middle value, this gives $d \approx 25$ μm. With $R = 2.5$ cm, we have $\beta = 10^{-3}$ (which is $\ll 1$). Thus a strength parameter in `pkdgrav` of 100 Pa is approximately 10^5 Pa in the frost layer bonding model.

Appendix C. Relating impact pressure and speed

Here we derive an order-of-magnitude estimate relating the impact speed between two bodies (v_{imp}) and the pressure exerted by that impact (P_{imp}). This relationship allows us to understand how the impact fragmentation limit (which is in units of speed) should scale with the bond strength (the bond's resistance to stress, which has units of pressure—see Section 2.1). After all, any stresses that break an aggregate should be self-consistent, be they impacts or tides.

We note that the following analysis is very similar to the technique employed by Hatzes et al. (1991) relating their experimentally determined critical collision speed for sticking to the critical force required for fragmentation in their experiments. This discussion goes one step further, attempting to convert from impact force to pressure by employing the results of Appendix B.

We begin with the simplified (one-dimensional) impulse equation relating the change in momentum of a rigid body to the force applied over a short time interval by an external body:

$$m\Delta v = F\Delta t \tag{46}$$

where m is the mass of the body impacted, Δv is the body's change in velocity due to the collision, F is the force applied during the collision (which is assumed to be constant), and Δt is the timescale of the collision—the time interval over which F is applied. This is a simplification of the force experienced during a collision, in which material compression exerts a varying pressure—we are merely aiming for an order-of-magnitude estimate.

Then we take the impact speed (v_{imp}) to be equal to Δv; this assumption introduces a worst case error of a factor of two, which depends on the coefficient of restitution of the impact.

Now we depart from Hatzes et al. (1991), and convert to pressure: the force of the impact is distributed over the contact area, or $P_{imp} = F/A$. To find A, we use the results of Appendix B, which were that the contact area between two equal-size spheres (of radius R) is $\pi R^2 \beta(1 - \beta/4)$, where β is the ratio of the particle overlap (or twice the deformation depth) to the particle radius.

With the mass of the body being $m = \frac{4}{3}\pi R^3 \rho$, we now have:

$$\frac{4\pi R^3 \rho v_{imp}}{3\pi R^2 \beta(1 - \beta/4)\Delta t} \sim \frac{R\rho v_{imp}}{\beta(1 - \beta/4)\Delta t} = P_{imp} \tag{47}$$

where the factor of $\frac{4}{3}$ has been neglected.

Putting this into units relevant for the rings of Saturn, and taking $\beta \ll 1$:

$$\left(\frac{1}{\beta}\right)\left(\frac{R}{1\text{ m}}\right)\left(\frac{\rho}{1\text{ g/cm}^3}\right)\left(\frac{v_{imp}}{1\text{ mm/s}}\right)\left(\frac{1\text{ s}}{\Delta t}\right) = \left(\frac{P_{imp}}{1\text{ Pa}}\right) \tag{48}$$

Now we apply this relationship to answer the original question: what impact pressure corresponds to a given impact speed?

For this example, we choose $v_{imp} = 0.5$ mm/s—approximately the escape speed of a sphere with $R = 1$ m, and $\rho = 0.5$ g/cm^3. (Note that this speed is the fragmentation limit used throughout our simulations.)

One of the larger uncertainties among the quantities in Eq. (48) is Δt. Hatzes et al. (1991) places the impact timescale at approximately 0.1 s, while Albers and Spahn (2006) have a timescale of ~ 1 ms (cf. their Fig. 2). For this exercise, we choose an intermediate value of 0.01 s, and recognize that we now have at least an order of magnitude of uncertainty. Last, as in Appendix B, we use $\beta = 10^{-3}$ (for the overlapping frost layer model).

These inputs yield a limiting impact pressure of 25,000 Pa, or about 10^4 Pa. Thus, in the overlapping frost layer cohesion model, and given the assumptions and simplifications made, an impact fragmentation limit of 0.5 mm/s is approximately consistent with an aggregate strength (tensile and shear) of 10^4 Pa.

If one were to instead use $\beta = 2$, as `pkdgrav` does (since it uses the entire particle cross-section as the contact area), one would instead find that the pressure is ~25 Pa. This is within an order of magnitude of 100 Pa—our choice of strength for our first suite—justifying our choice for that parameter in our simulations.

References

Albers, N., Spahn, F., 2006. The influence of particle adhesion on the stability of agglomerates in Saturn's ring. Icarus 181, 292–301.

Barnes, R., Quinn, T.R., Lissauer, J.J., Richardson, D.C., 2009. *N*-body simulations of growth from 1 km planetesimals at 0.4 AU. Icarus 203, 626–643.

Borderies, N., Goldreich, P., Tremaine, S., 1984. Unsolved problems in planetary ring dynamics. In: Greenberg, R., Brahic, A. (Eds.), Planetary Rings. Univ. of Arizona Press, Tucson, pp. 713–734.

Bridges, F.G., Supulver, K.D., Lin, D.N.C., Knight, R., Zafra, M., 1996. Energy loss and sticking mechanisms in particle aggregation in planetesimal formation. Icarus 23, 422–435.

Bridges, F., Supulver, K., Lin, D.N.C., 2001. Energy loss and aggregation processes in low speed collisions of ice particles coated with frosts or methanol/water mixtures. In: Poschel, T., Luding, S. (Eds.), Granular Gases. Springer, Berlin/Heidelberg, pp. 153–183.

Colwell, J.E., Esposito, L.W., Sremcevic, M., Stewart, G.R., McClintock, W.E., 2007. Self-gravity wakes and radial structure of Saturn's B ring. Icarus 190, 127–144.

Cuzzi, J., Clark, R., Filacchione, G., French, R., Johnson, R., Marouf, E., Spilker, L., 2009. Ring particle composition and size distribution. In: Dougherty, M.K., Esposito, L.W., Krimigis, S.M. (Eds.), Saturn from Cassini–Huygens. Springer, Dordrecht, pp. 459–510.

Daisaka, H., Ida, S., 1999. Spatial structure and coherent motion in dense planetary rings induced by self-gravitational instability. Earth Planets Space 51, 1195–1213.

French, R.G., Salo, H., McGhee, C., Dones, L., 2007. HST observations of azimuthal asymmetry in Saturn's rings. Icarus 189, 493–522.

Hatzes, A.P., Bridges, F., Lin, D.C., Sachtjen, S., 1991. Coagulation of particles in Saturn's rings—Measurements of the cohesive force of water frost. Icarus 89, 113–121.

Hill, G.W., 1878. Researches in lunar theory. Am. J. Math. 1, 5–26.

Johnson, K.L., Kendall, K., Roberts, A.D., 1971. Surface energy and the contact of elastic solids. Proc. R. Soc. Lond. A. 324, 310–313.

Karjalainen, R., 2007. Aggregate impacts in Saturn's rings. Icarus 189, 523–537.

Karjalainen, R., Salo, H., 2004. Gravitational accretion of particles in Saturn's rings. Icarus 172, 328–348.
Lewis, M.C., Stewart, G.R., 2000. Collisional dynamics of perturbed planetary rings. I. Astron. J. 120, 3295.
Lois, G., Carlson, J.M., 2007. Force networks and the dynamic approach to jamming in sheared granular media. Europhys. Lett. 80, 58001.
Marouf, E.A., Tyler, G.L., Zebker, H.A., Simpson, R.A., Eshleman, V.R., 1983. Particle size distributions in Saturn's rings from Voyager 1 Radio Occultation. Icarus 54, 189–211.
Porco, C.C., Weiss, J.W., Richardson, D.C., Dones, L., Quinn, T., Throop, H., 2008. Simulations of the dynamical and light-scattering behavior of Saturn's rings and the derivation of ring particles and disk properties. Astron. J. 136, 2172.
Quinn, T., Perrine, R.P., Richardson, D.C., Barnes, R., 2010. A symplectic integrator for Hill's equations. Astron. J. 139, 803–807.
Richardson, D.C., 1995. A self-consistent numerical treatment of fractal aggregate dynamics. Icarus 115, 320–335.
Richardson, D.C., Quinn, T., Stadel, J., Lake, G., 2000. Direct large-scale *N*-body simulations of planetesimal dynamics. Icarus 143, 45–59.
Richardson, D.C., Michel, P., Walsh, K.J., Flynn, K.W., 2009. Numerical simulations of asteroids modelled as gravitational aggregates with cohesion. Planet. Space Sci. 57, 183–192.
Robbins, S.J., Stewart, G.R., Lewis, M.C., Colwell, J.E., Sremcevic, M., 2010. Estimating the masses of Saturn's A and B rings from high-optical depth *N*-body simulations and stellar occultations. Icarus 206, 431–445.
Salo, H., 1992. Gravitational wakes in Saturn's rings. Icarus 359, 619–621.
Salo, H., Karjalainen, R., French, R.G., 2004. Photometric modeling of Saturn's rings. II. Azimuthal asymmetry in reflected and transmitted light. Icarus 170, 70–90.
Scheeres, D.J., 1998. The restricted Hill four-body problem with applications to the Earth–Moon–Sun system. Celest. Mech. Dynam. Astron. 70, 75–98.
Spahn, F., Albers, N., Sremcevic, M., Thornton, C., 2004. Kinetic description of coagulation and fragmentation in dilute granular particle ensembles. Europhys. Lett. 67, 545–551.
Spitale, J.N., Porco, C.C., 2010. Free unstable modes and massive bodies in Saturn's outer B ring. Astron. J. 140, 1747–1757.
Stadel, J., 2001. Cosmological *N*-body simulations and their analysis. Thesis, University of Washington, Seattle, 126pp.
Supulver, K.D., Bridges, F.G., Tiscareno, S., Lievore, J., Lin, D.N.C., 1997. The sticking properties of water frost produced under various ambient conditions. Icarus 129, 539–554.
Tiscareno, M.S., Perrine, R.P., Richardson, D.C., Hedman, M.M., Weiss, J.W., Porco, C.C., Burns, J.A., 2010. An analytic parameterization of self-gravity wakes in Saturn's rings, with application to occultations and propellers. Astron. J. 139, 492–503.
Tremaine, S., 2003. On the origin of irregular structure in Saturn's rings. Astron. J. 125, 894–901.
Wisdom, J., Tremaine, S., 1988. Local simulations of planetary rings. Astron. J. 95, 925–940.
Zebker, H.A., Marouf, E.A., Tyler, G.L., 1985. Saturn's rings: Particle size distributions for thin layer model. Icarus 64, 531–548.