

Abstract

Title of Dissertation: *N*-body Simulations with Cohesion in Dense Planetary Rings

Randall Patrick Perrine, Doctor of Philosophy, 2011

Dissertation directed by: Professor Derek C. Richardson
Department of Astronomy

This dissertation is primarily focused on exploring whether weak cohesion among icy particles in Saturn’s dense rings is consistent with observations—and if so, what limits can be placed on the strength of such cohesive bonds, and what dynamical or observable consequences might arise out of cohesive bonding.

Here I present my numerical method that allows for *N*-body particle sticking within a local rotating frame (“patch”)—an approach capable of modeling hundreds of thousands or more colliding bodies. Impacting particles can stick to form non-deformable but breakable aggregates that obey equations of rigid body motion.

I then apply the method to Saturn’s icy rings, for which laboratory experiments suggest that interpenetration of thin, frost-coated surface layers may lead to weak bonding if the bodies impact at low speeds—speeds that happen to be characteristic of the rings. This investigation is further motivated by observations of structure in the rings that could be formed through bottom-up aggregations of particles (i.e., “propellers” in the A ring, and large-scale radial structure in the B ring).

This work presents the implementation of the model, as well as results from a suite of 100 simulations that investigate the effects of five parameters on the equilibrium characteristics of the rings: speed-based merge and fragmentation limits, bond strength, ring surface density, and patch orbital distance (specifically the center of either the A or B ring), some with both monodisperse and polydisperse particle comparison cases.

I conclude that the presence of weak cohesion is consistent with observations of the A and B rings, and present a range of parameters that reproduce the observed size distribution and maximum particle size. The parameters that match observations differ between the A and B rings, and I discuss the potential implications of this result. I also comment on other observable consequences of cohesion for the rings, such as optical depth and scale height effects, and discuss the unlikelihood that very large objects are grown bottom-up from cohesion of smaller ring particles.

Lastly, I include a brief summary of other projects in ring dynamics I have undertaken before and during my thesis work.

***N*-body Simulations with Cohesion in Dense Planetary Rings**

by

Randall Patrick Perrine

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland at College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2011

Advisory Committee:

Professor Derek C. Richardson, Chair

Professor M. Coleman Miller

Professor Douglas Hamilton

Professor Jessica Sunshine

Professor Richard J. Walker

Dr. Matthew Tiscareno

© Randall Patrick Perrine 2011

Preface

Material in Chapters 1 to 4 and 6, as well as Appendix A, have appeared in published work (Perrine et al. 2011, *Icarus*, vol. 212, pp. 719–735) and work submitted for publication (Perrine and Richardson 2011, *Icarus*, submitted). Material throughout this dissertation has been presented at meetings of the Division for Planetary Sciences of the American Astronomical Society via abstracts, oral presentations, and posters. Section 5.2 refers to work published in Tiscareno et al. (2010a), and Section 5.3 to work published in Quinn et al. (2010).

This dissertation is the result of the efforts of myself as well as others, and—where possible—the contributions of others are highlighted in this dissertation. Essential to this project was the adding of new functionality to existing code, extensive debugging and testing of this code, and then carrying out a large number of supercomputer and cluster-based simulations. Prof. Richardson assisted with the writing of the code while we added the new functionality, and provided general advising, but all testing, revising, running, and analyzing that followed is my own work. A coauthor, Prof. Daniel J. Scheeres (U Colorado), derived equations of motion for the solid-body aggregate dynamics, which appeared in Perrine et al. (2011).

For my parents, and my loving wife.

Acknowledgements

I would like to thank all of the graduate students, faculty, and staff at the Astronomy Department at the University of Maryland, College Park, for fostering an encouraging and enlightening environment in which I could learn and work. Most especially, I thank my advisor, Derek, who made my journey a relaxing adventure. I also thank the CTC, Department of Astronomy, and OIT computing staffs who maintained the `yorp` and `deephought` computer clusters, without whom I could not possibly have completed this 540,000+ CPU hour (61+ CPU year!) project.

Contents

List of Tables	vii
List of Figures	viii
1 Introduction	1
1.1 Background: Cohesion in Saturn’s Rings?	1
1.2 Background: Modeling Planetary Rings	4
1.3 Terminology	9
1.4 Overview of This Dissertation	12
2 Methods for Cohesion Simulations	13
2.1 pkdgrav	13
2.2 The Aggregate Model	15
2.3 Code Alterations	21
2.3.1 Three Coordinate Systems	23
2.3.2 Aggregates and Periodic Boundaries	24
2.3.3 Aggregate Orientation	27
2.3.4 Aggregate Acceleration, Torque, and Stress	29
2.3.5 Collisions	30
2.3.6 Overlapping Particles and Correction Strategies	32
3 Cohesion Simulations	35
3.1 Parameters	35
3.2 Data Analysis	45
4 Results and Discussion of Cohesion Simulations	49
4.1 The Lifetime of a Typical Aggregate	49
4.2 Effects of Surface Density	51
4.3 Dominant Fragmentation Mechanism: Collisions or Stress?	53
4.4 Constraints on A Ring Bonding Parameters	56
4.5 Constraints on B Ring Bonding Parameters	60
4.6 Reconciling Observation and Model Results	62

4.7	Maximum Aggregate Size in Equilibrium	65
4.8	Other Observational Diagnostics for Constraining Aggregation Parameters	68
4.9	Validity of Monodisperse Starting Conditions	73
4.10	Comments on the Range of Bonding Parameters	74
5	Other Rings Work	76
5.1	Review of Work in Ring-Moonlet Interactions	76
5.1.1	Model Alterations	78
5.1.2	Early Results: Edgewaves and Gap-clearing	81
5.2	An Analytical Parameterization of Wake/Interwake Optical Depth Contrast	88
5.3	A Symplectic Method for Local Simulations	90
6	Summary and Conclusions	91
6.1	Summary	91
6.2	Conclusions	93
A	Derivations	96
A.1	Contact Area for Intersecting Spheres	96
A.2	A Simple Model Relating Impact Pressure and Speed	98
	Bibliography	101

List of Tables

3.1	Table of Simulation Parameters: A Ring (Part 1)	41
3.2	Table of Simulation Parameters: A Ring (Part 2)	42
3.3	Table of Simulation Parameters: B Ring (Part 1)	43
3.4	Table of Simulation Parameters: B Ring (Part 2) and Orbital Distance . .	44

List of Figures

1.1	Generic Diagram for Local (Sliding Patch) Simulations	6
1.2	Example Snapshot of a Local (Sliding Patch) Simulation	7
2.1	Small-scale Test of the Coherent Aggregate Model	16
2.2	Schematic Diagram: Aggregation with Periodic Boundaries	22
3.1	Two Example Aggregate Size Distributions	46
3.2	Temporal Variation of Aggregate Population Properties	47
4.1	Snapshots of Two A Ring Simulations	50
4.2	Snapshots of Two B Ring Simulations	51
4.3	Aggregate Population Properties vs. Σ	52
4.4	Aggregate Population Properties vs. Bond Strength	54
4.5	Aggregate Population Properties vs. Merge Limit (1 v_{esc})	57
4.6	Aggregate Population Properties vs. Merge Limit (2 v_{esc})	58
4.7	Aggregate Population Properties vs. Fragmentation Limit	59
4.8	Optical Depth (τ_{phys}) Results for All Simulations	69
4.9	Scale Height vs. Orbital Semimajor Axis	70
4.10	Scale Height vs. τ_{phys}, R_{max}	71
5.1	Overview of Gaps in the Outer A Ring	77
5.2	Schematic: Edgewave Theory	81
5.3	Demonstration of a Gap-Moon Interaction: Edgewaves	82
5.4	Demonstration of a Gap-Moon Interaction: Gap Width Variation with Moonlet Mass	83
5.5	Results of the Gap-Moon Interaction Study: Gap Width and Edgewave Wavelength Variation with Moonlet Mass	84
5.6	Demonstration of a Gap-Moon Interaction: An Eccentric Moonlet	87
A.1	Diagram for the Simple Contact Area Model	97

Chapter 1

Introduction

1.1 Background: Cohesion in Saturn's Rings?

The rings of Saturn are an immensely wide and flat shearing disk made up of icy particles ranging in size approximately from 1 cm to 20 meters. The main rings extend about from 67,000 km (the inner edge of the D ring) to 137,000 km from Saturn (the outer edge of the A ring). As large as the rings are, however, their vertical scale is quite small—in equilibrium, the rings are only as thick as a few times the size of the largest particles in the rings. Thus the rings are about 250,000 km wide, but only about 10 meters thick (Greenberg and Brahic 1984).

The two most massive rings, the A and B rings, are the densest rings in the Solar System, making them ideal laboratories to study the dynamics of self-gravitating and collisionally evolving systems. Direct collisions between particles are an important dynamical process in the A and B rings, as collisions dissipate energy. Using particle density and velocity dispersion estimates, a particle collides approximately 2–3 times per orbital

period in these dense rings.¹

As the small particles comprising the rings of Saturn are composed almost entirely of water ice (see Cuzzi et al. 2009, for a recent review), many experimentalists have undertaken the task of measuring the coefficient of restitution² in icy particle collisions (e.g., Bridges et al. 1984; Supulver et al. 1995). Interestingly, some of these experiments have revealed that frost-coated ice can stick at low impact speeds, forming a cohesive bond (Hatzes et al. 1991; Bridges et al. 1996, 2001; Supulver et al. 1997). In particular, Hatzes et al. (1991) describe this cohesion by invoking a “Velcro” model—the rough texture of the frost layers provide an interlocking structure for colliding particles. Particularly intriguing is their discovery of a critical impact speed for cohesion (hereafter called a “merge limit”) of ~ 0.3 mm/s (for 2.5 cm-radius bodies), as this speed is on the order of the gravitational escape speed (for a test particle on the surface) of an icy sphere 1 m in radius. If this cohesion mechanism is generic across size scales, and the particles in the dense rings of Saturn are experiencing cohesion, they may form large aggregations of material.

Interestingly, analysis of recent observations of Saturn’s rings implies the need for such a bonding mechanism. Porco et al. (2008) show that the contrast in the azimuthal brightness asymmetry³ 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. Tremain (2003) suggests that these zones occupy a region of

¹A simple lower limit comes from particles crossing the ring plane, plunging through ring material with optical depths near unity, on inclined orbits twice per orbit.

²The ratio of the rebound to impact speeds in the direction along the line of centers of the colliding particles.

³This asymmetry is a variation in reflected and transmitted light with incidence angle in Saturn’s dense rings.

dynamical phase-space that is “shear-free”—zones with zero Keplerian shear—and thus subject to a solid-liquid phase transition. The theory is that these zones may consist of large-scale structures, consisting of 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.

In addition, observations of the rings (Marouf et al. 1983 and Zebker et al. 1985 with Voyager radio occultation experiments, and French and Nicholson 2000 with stellar occultations; see Cuzzi et al. 2009 for a summary) have revealed indications of the size distribution and maximum particle size within different regions of the rings. Overall, the results of their modeling and data analysis are that the size distribution and maximum particle size in the A and B rings are roughly independent of the distance from Saturn. Adopting a functional form for the size distribution of $n(R) \propto R^\alpha$ (where $n(R)$ is the number of particles of radius R , and α is a parameter to be determined), they find α to be roughly between -2.5 and -3, and the maximum particle size⁴ (R_{max}) to be 5–20 m (with the Voyager results favoring 5–10 m, and the stellar occultation showing a nearly constant 20 m). There may also be trends with orbital distance from Saturn: each experiment showed a steepening of the size distribution from the inner to the outer A ring (α parameter from -2.7 to -3 for Voyager, and from -2.75 to -2.9 for the stellar occultations), and there are hints that R_{max} may increase with distance from Saturn (e.g., the Voyager experiment showed R_{max} increasing from 5–6 m in the inner and mid A ring to 9–11 m in the Encke gap region; see Cuzzi et al. 2009). Note that these experiments are subject to different modeling assumptions (e.g., the minimum particle size, or a “classical” versus “finite thickness” ring geometry), and may be somewhat skewed by structure that may

⁴Note that R_{max} may describe a single large body that, due to its low number frequency, does not necessarily lie on the overall size distribution. Its size could be well above the predicted cutoff.

vary with Saturnian distance (i.e., wakes; see Section 1.3).

These findings form the primary motivation for this project, and have elicited the following questions that this dissertation will attempt to answer: Can cohesive aggregates form in a dense planetary ring environment? Do they disrupt known ring structure and dynamics? What parameters allow cohesive aggregates to form while remaining consistent with existing observations, and are those parameters reasonable? Are there other consequences of aggregation that can be observed to further constrain this process? And finally, can the presence of cohesive aggregates explain the existence of unexplained structure in the rings? Answering these questions should also provide insight into the yet-unknown surface properties of the particles comprising the rings.

1.2 Background: Modeling Planetary Rings

Studying the full effect of cohesive bonding 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, planetary tides, and interparticle collisions. It may prove difficult to further incorporate dynamical bonding and fragmentation in a fully self-consistent way into current analytical models. 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), and a few groups study planetary ring dynamics with a large number of bodies via local N -body simulations (e.g., Lewis and Stewart 2000; Karjalainen and Salo et al. 2004; Porco et al. 2008), but none combine these models to study the emergent behavior of tens of thousands (or even millions) of N -body particles interacting with cohesion over many orbital times. So, for this project, I created a new model that incorporates cohesion among N -body particles into a local, rotating frame—discussed in detail in Chapter 2.

Local simulations are valuable tools for these studies: by restricting the computational volume to a small region of interest (called 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 meter 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 with periodic boundary conditions (also called a “sliding patch” model), is a computational shortcut that was first used by Julian and Toomre (1966), and applied to dense planetary rings by Wisdom and Tremaine (1988). It can be applied to many non-inertial systems orbiting a large body, such as debris disks around stars. Local simulations use Hill’s linearized equations of motion (Hill 1878; Wisdom and Tremaine 1988):

$$\begin{aligned}
 \ddot{x} &= \mathcal{F}_x + 3\Omega^2 x + 2\Omega \dot{y}, \\
 \ddot{y} &= \mathcal{F}_y - 2\Omega \dot{x}, \\
 \ddot{z} &= \mathcal{F}_z - \Omega^2 z.
 \end{aligned}
 \tag{1.1}$$

with \mathcal{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. 1.1.

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. Par-

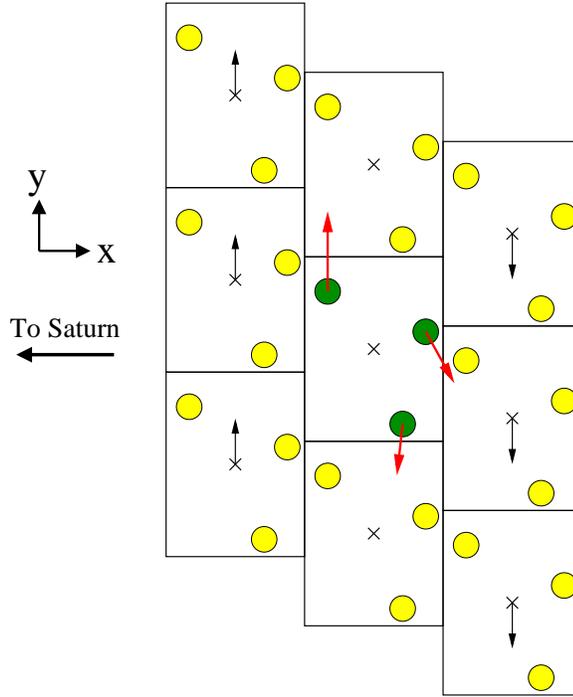


Figure 1.1: A schematic diagram of a local sliding patch model with shearing periodic boundary conditions. The three green particles in the center box are the simulated particles; yellow particles are in replicated patches that provide boundary conditions. The x coordinate is the radial direction, with 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 black 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 a boundary of the central patch 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 $+\frac{3}{2}\Omega L_x \hat{y}$ (see section 1.2). Note that while only one ring of replicated patches is shown, it is customary to use three rings of replicas, in order to provide a smooth background.

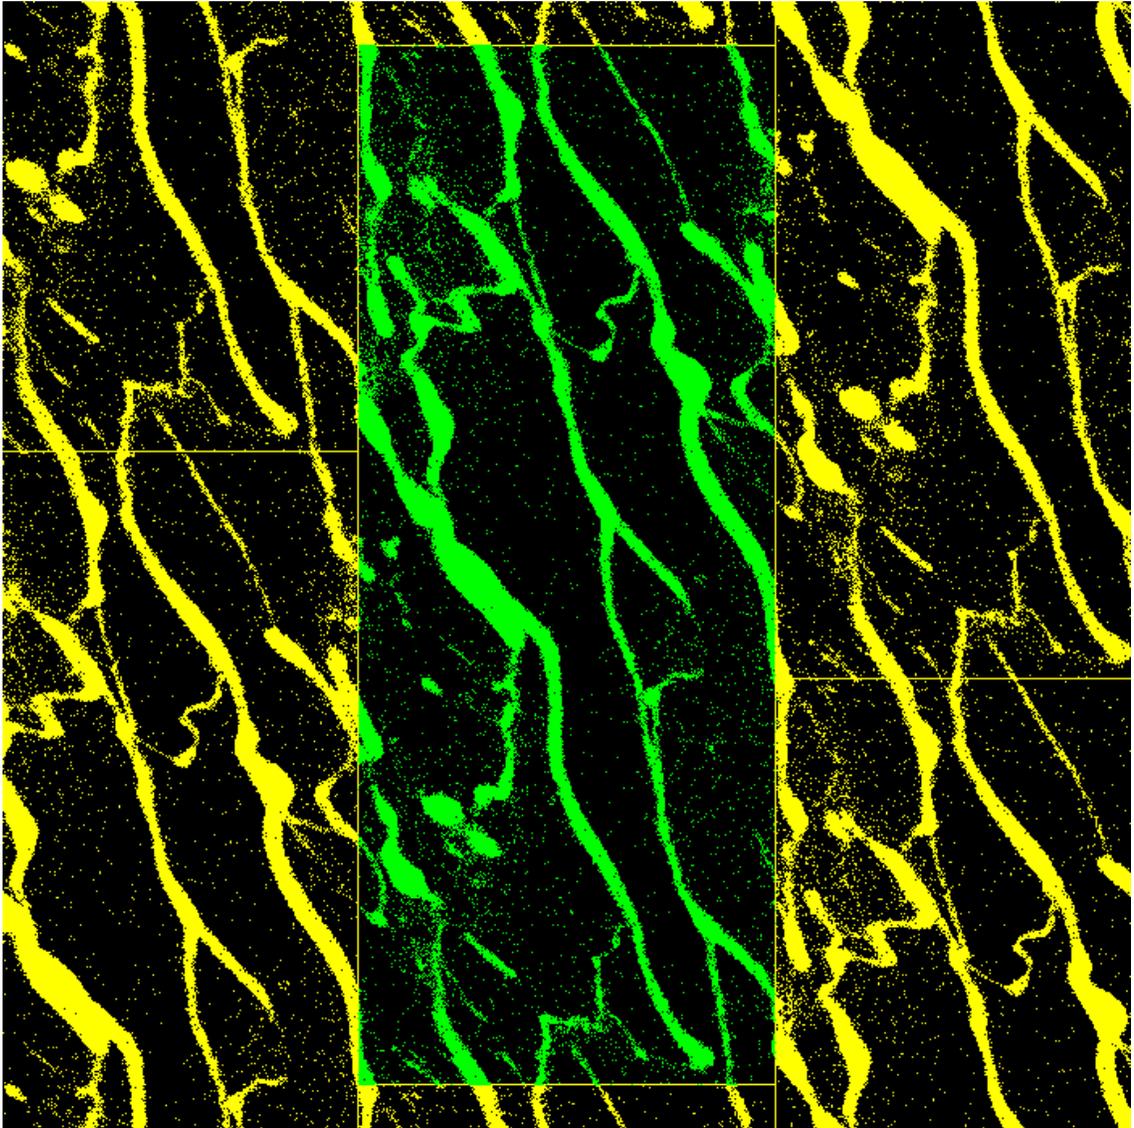


Figure 1.2: Snapshot from a local simulation with periodic boundary conditions. Viewing geometry is as in Fig. 1.1. Particles in green are the $\sim 86,000$ simulated bodies, and the yellow particles replicas of the greens to provide boundary conditions. This simulation is based in the outer A ring, with a surface density of 500 kg/m^2 , using the same parameters as a fiducial A ring simulation outlined in Section 3.1, but without any interparticle cohesion. Self-gravity and collisions are enabled, as usual, which is evident in the ready formation of the prevalent wake structure (see Section 1.3), with scattered individual particles in the interwake regions.

ticles 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. I 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. 1.1.

It is important to keep in mind that these equations of motion are linearized, which is a valid approximation so long as the particles’ positions with respect to the patch center ($r \sim \sqrt{x^2 + y^2}$) are much smaller than the semimajor axis of the patch’s orbit (a); that is, distortion induced by the linearization will be apparent as r^2/a^2 approaches unity. Thus the dimensions of the patch (in the azimuthal and radial directions) are kept much smaller than the patch’s semimajor axis.

Many groups employ this local “patch” model to study planetary rings (recent examples include Lewis and Stewart 2000; Karjalainen and Salo et al. 2004; Porco et al. 2008), though none of those groups employ interparticle bonding in their models. Thus my 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) include 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. My new method allows for studies into more realistic shape and spin distributions via agglomeration, and accounts for (admittedly simplified) breakage due to collisions and external forces.

It is important to note that while this new model does not attempt to encapsulate all of the complex physics relevant to granular mechanics, it remains more sophisticated than previous methods. In applying my model to the rings of Saturn, I 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.

1.3 Terminology

In my model, cohesion leads to the formation of “aggregates” that are collections of particles joined 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 an aggregate remain fixed with respect to one another as the aggregate moves and rotates in response to external forces and torques. Aggregates can agglomerate through collisional accretion with other bodies—free particles, or other aggregates—and can fragment through collisional events or stress fragmentation (when stresses exceed bond strengths).

Other researchers (e.g., Karjalainen and Salo 2004, and Karjalainen 2007) use the term “aggregate” to refer to a collection of particles bound together by gravity alone. I refer to these structures as “gravitational aggregates.” Unless indicated otherwise, in this project, “aggregate” refers to a bonded aggregate, though I use “bonded aggregate” at times for clarity. Both types of structures appear in my simulations (cf. Section 4.1, Figs.

1.2, 4.1, and 4.2).

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 size estimated by twice the critical wavelength of the region (the shortest axisymmetric wavelength stabilized by differential rotation alone): $\lambda_{crit} = 4\pi^2 G\Sigma/\kappa^2$ (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 degrees, depending on local conditions.

Researchers (e.g., Salo 2004, and 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 my simulations (see, for example, Figs. 1.2, 4.1, and 4.2).

Important characteristic speed scales in my simulations include the particle escape speed (v_{esc}) and shear speed (v_{shear}). The former has the usual definition: the minimum initial speed of a massless test particle needed to reach infinity relative to an isolated ring particle, $v_{esc} = \sqrt{\frac{2Gm}{R}}$, where m and R are the mass and radius of the ring particle. For this definition, I ignore planetary gravity, tidal effects, and perturbations from other

particles (or moons). The shear speed is the relative speed between two (massless) objects on circular orbits with semimajor axes separated by the sum of their radii, that just barely touch as the inner body overtakes the outer; in the linearized approximation (see Wisdom and Tremaine 1988), $v_{shear} = \frac{3}{2}\Omega\delta a$, where δa is the sum of the particles' radii. The bodies are massless in this definition so that their mutual gravitational attraction does not increase their relative speed. The escape and shear speeds are relevant impact speed scales for particles in dense rings.

Finally, I define a quantity β that specifies the contact area assumed when discussing different bonding models, while at the same time encapsulating complex bonding physics into a simple term (similar to the concept of a coefficient of restitution). When any model calculates bond strength from relative acceleration, that model must specify the assumed contact area, as the contact area is the conversion factor between force and pressure. Thus, in this work, I quote the assumed β whenever a bond strength is discussed.

For the full derivation and discussion of β , I refer the reader to Appendix A. In short, $\beta \equiv d/R$, or the ratio between the depth of particle surface layer interpenetration (d) to the radius of the equal-size bodies (R). (For simplicity, we restrict ourselves here to the monodisperse case.) For perfectly overlapping spheres, $\beta = 2$, and the resulting contact area is the cross-section of the particles, πR^2 (this is the contact area typically assumed in the model). If $\beta = 0$, then there is no overlap, and the contact area is zero. At small values of β , the contact area approximately equals the particles' cross-section times β .

Contact area scales inversely with strength (with constant force). As such, β provides a simple means to convert from any strength used as a parameter in the code to a strength meaningful in whatever bonding model is under study: simply divide the parameter by β (or $\beta(1 - \beta/4)$ if $\beta \sim 1$). For example, the bonding concept used in this work is based on overlapping frost layers; Hatzes et al. (1991) studied 2.5 cm ice spheres, and I estimate $d \approx 25 \mu\text{m}$ from their results. With $R = 2.5$ cm, we have $\beta = 10^{-3}$ (which is $\ll 1$). Thus

a strength parameter of 100 Pa in the code is approximately 10^5 Pa in the frost layer bonding model.

1.4 Overview of This Dissertation

Chapters 2 through 4 focus on the cohesion model project. Chapter 2 discusses the cohesion model itself in detail, including considerations and alterations to the code needed to allow for cohesion among particles in a local simulation. Chapter 3 discusses the cohesion simulations, outlines and justifies the parameters used therein, and discusses how the simulation data was analyzed. Chapter 4 presents and discusses the results from these simulations, and addresses the questions at the end of Section 1.1. Chapter 5 details other work accomplished while exploring possible topics for a thesis on planetary ring dynamics, and as side projects undertaken during the cohesive ring particle project: Section 5.1 examines work that was never submitted for publication, but could form the basis for future students interested in planetary ring dynamics, and Sections 5.2 and 5.3 discuss published work. Finally, Chapter 6 provides an overall summary and conclusions.

Chapter 2

Methods for Cohesion Simulations

This chapter is organized as follows: first, I present the general structure of the code `pkdgrav`, which (with the exception of the new symplectic method for local simulations) was largely in place prior to my involvement in this project. Next, I describe the mechanics of the aggregate model, explain how the particles form into and break away from aggregates, and present the equations of motion for aggregates. The last section details the changes I made to the existing code in order to implement the aggregation model in a local rotating frame.

2.1 `pkdgrav`

The aggregation model is built as an optional feature 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 (see, e.g., Barnes and Hut 1986 and Richardson 1994 for a discussion of tree code concepts). 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. The model does this using a neighbor search borrowed from Smoothed-Particle Hydrodynamic (SPH) techniques (for a general review of SPH, see Monaghan 1992). In this method, each particle searches for its nearest n_{smooth} neighbors and calculates whether its path would intersect with any of these neighbors during the next timestep. If so, the corresponding collisions are carried out, and the positions and velocities of the particles are updated as needed. (For more discussion, see 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.¹ Quinn et al. (2010) present a symplectic leapfrog method for the rotating patch frame that I have implemented into the code. One complication of the method is the introduction of a new canonical variable, the momentum P_y ,² that needs to be tracked for each particle. (This quantity is conserved in collisions, and is treated in a similar manner to a body’s angular momentum.) Impulsive events, like collisions, require updates to P_y during the drift phase. I adapted Eqs. (33) and (34) of Quinn et al. (2010) to find P_y during the drift:

$$P_y = \dot{y} + 2\Omega [x_{event} + \dot{x}(\Delta t/2 - t_{event})] \quad (2.1)$$

where \dot{x} and \dot{y} are the (constant) radial and azimuthal velocity components during the drift, Ω is the usual orbital frequency of the rotating coordinate system, x_{event} is the (drifted)

¹This work is also discussed in Section 5.3.

² $P_y \equiv \dot{y} + 2\Omega x$; see text for details.

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). The main issue is that P_y is defined and calculated for all particles *before* the opening kick, but the drift phase adjusts a particle’s position using its velocity *after* the opening kick. Therefore, in order to properly compute P_y at any arbitrary time during the drift phase, the code 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. (2.1) 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. (2.1) 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)).

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.

2.2 The Aggregate Model

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. I 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.”

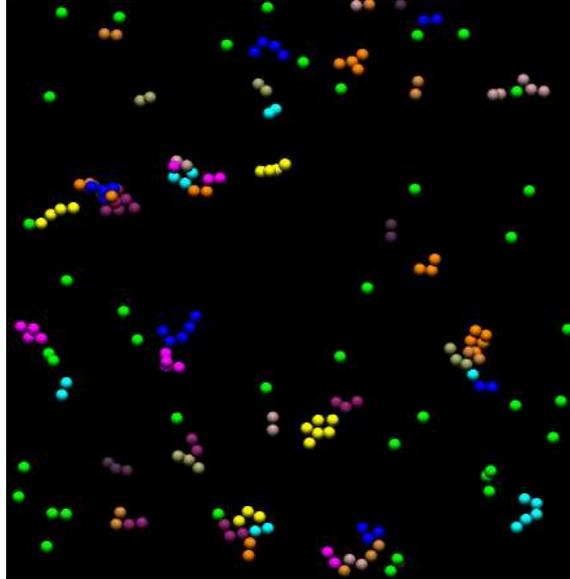


Figure 2.1: Snapshot from an early small-scale test simulation containing 200 particles in the cohesive particle sliding patch model, with merging and fragmentation enabled. The particles are all 1 m in radius. Green bodies are free (unbonded) particles; when two free particles merge, they become an 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. 1.1. Note that the appearance of major particle overlap is a projection effect, as this code is a 3D model.

(This threshold 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—the option for a single fixed value is assumed throughout the remainder of this dissertation). Fractal growth proceeds as more free particles—or other aggregates—collide and stick in this way. In my 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 the model includes two breaking mechanisms that each 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.” (This parameter can also be either an absolute speed or a value scaled by the mutual escape speed of the colliding bodies—again, the fixed value option is assumed throughout the remainder of this work). 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 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,³ 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. Indeed, 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 interpen-

³The strength can be made size-dependent according to $S \propto R^\gamma$, where R is the radius of the aggregate’s minimally enclosing sphere, and γ is a user-defined constant. In this context, γ is assumed to be zero throughout all simulations discussed in Chapter 3.

etration of the surfaces of the bodies) will have a drastically different contact area. This strategy 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 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 by the β parameter (see Section 1.3 and Appendix A.1).

Note that there is no internal consistency check to verify that the user’s choice of fragmentation limit matches physically with 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 (cf. Section 4.3). So it is up to the user to match these two parameters physically, if desired. In order to assist with this, I 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 A.2). This relationship includes the generic parameter β (defined in Appendix A.1) that scales the strength by considering the varying contact areas of different bonding methods.

Also note that the code does 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, the code 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, these 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 this cohesion model is only an approximation to the behavior of real cohesive materials, but I believe the code captures the essential elements while keeping the computations tractable.

However, the model’s generality remains an advantage, as it captures the basic behavior of bonding without limiting the range of cohesion mechanisms that it can mimic. For example, the 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.

Aggregates obey Euler’s equations of rigid body rotation:

$$\begin{aligned}
 I_1 \dot{\omega}_1 - \omega_2 \omega_3 (I_2 - I_3) &= N_1 \\
 I_2 \dot{\omega}_2 - \omega_3 \omega_1 (I_3 - I_1) &= N_2 \\
 I_3 \dot{\omega}_3 - \omega_1 \omega_2 (I_1 - I_2) &= N_3
 \end{aligned}
 \tag{2.2}$$

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:

$$\begin{aligned}
 \dot{\hat{\mathbf{p}}}_1 &= \omega_3 \hat{\mathbf{p}}_2 - \omega_2 \hat{\mathbf{p}}_3 \\
 \dot{\hat{\mathbf{p}}}_2 &= \omega_1 \hat{\mathbf{p}}_3 - \omega_3 \hat{\mathbf{p}}_1 \\
 \dot{\hat{\mathbf{p}}}_3 &= \omega_2 \hat{\mathbf{p}}_1 - \omega_1 \hat{\mathbf{p}}_2
 \end{aligned}
 \tag{2.3}$$

where $\hat{\mathbf{p}}_i$ denote the principal axes. Note that these are optimized for use in an aggregate’s

body frame, in which the inertia tensor is diagonalized. (This further simplifies the computations 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 $(\mathbf{r}_i - \mathbf{r}_a) \times (\dot{\mathbf{r}}_i - \dot{\mathbf{r}}_a)$ —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.). The code uses 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. (2.2) and (2.3). The aggregate centers of mass are advanced according to Eq. (1.1) 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 the 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 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 2.3.6 for a detailed discussion of particle overlaps.

2.3 Code Alterations

Two essential issues required direct attention while incorporating the rigid aggregate model of Richardson et al. (2009) into the local rotating frame model, which I outline here, with details reserved for the separate subsections that follow.

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. 2.2)—complicating, for example, collision outcome handling, and calculation of the properties of the aggregate’s center of mass (COM). To address this, the model stores 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. 2.2 intersects with a boundary), the code stores not only the standard in-patch position (shown in green in Fig. 2.2), 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). I 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

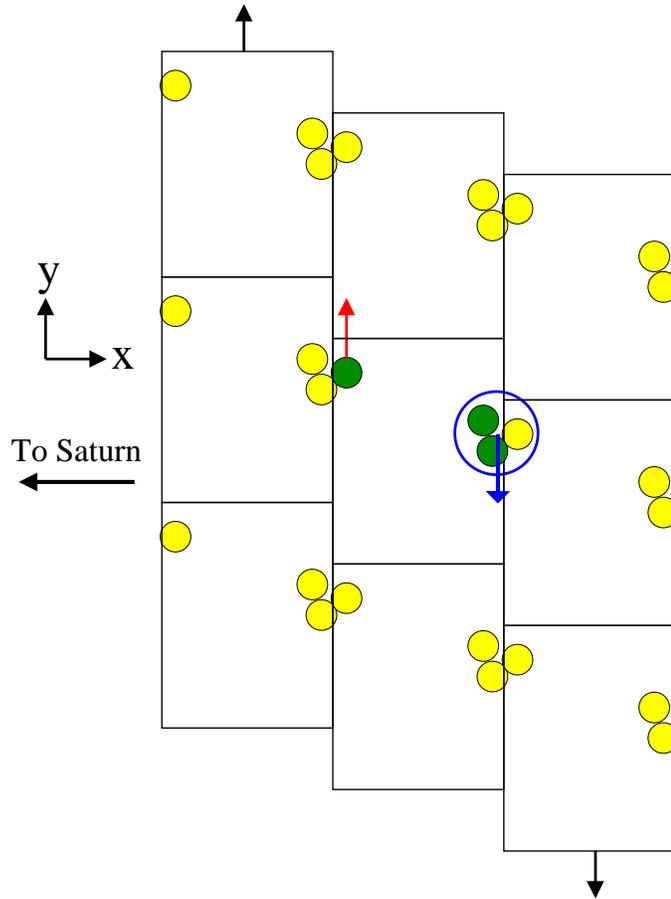


Figure 2.2: A schematic diagram similar to Fig. 1.1 illustrating the considerations involved with handling aggregates in a periodic patch. There are three particles in this system, and they are joined into a single aggregate. The aggregate is encircled, and the center of mass (COM) of this aggregate would lie at the approximate center of the circle (as drawn). The green 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 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. 1.1, 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 naively used green particle velocities, it would include the contribution of the red arrow, which would produce an incorrect result. Thus, the code 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).

effects on such processes as collision detection and resolution, and requires rederiving the equations of motion for an aggregate in a rotating frame; these equations were provided by coauthor Scheeres of Perrine et al. (2011), and appear as Appendix A in that work.

The remainder of this section describes explicitly all of the modifications and additions made to the numerical method to place the rigid aggregate model into the rotating local frame.

2.3.1 Three Coordinate Systems

To begin, I present the three sets of Cartesian coordinate systems (“frames”) in this model, then discuss their function, and finally describe how one transforms 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 used the most often, as the sliding patch model is based in this non-inertial frame (see Fig. 1.1). Its axes are initially aligned with the space frame at time $t = 0$; but as the patch orbits, it rotates about its \hat{z} axis at Ω , the orbital frequency of the patch.⁴ Translational motions for all bodies are calculated in the patch frame, using Hill’s equations of motion (Eq. (1.1)), which include the coriolis and centrifugal accelerations.

⁴The 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 .

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 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 Λ^{-1} , i.e. the inverse of the rotation matrix. To properly handle particle spin, first add $\Omega\hat{z}$ to the spin (to put the spin into the space frame), and then multiply by Λ^{-1} .

2.3.2 Aggregates and Periodic Boundaries

As mentioned at the top of Section 2.3, 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. 2.2). 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), I simply track the aggregate's constituent particles more completely. If it is joined

to an aggregate, I utilize two sets of coordinates to store the position of a single particle: the “wrapped patch position” (position #1), and the “unwrapped patch position” (position #2).

Position #1 is standard in the sliding patch model: the position of the particle in patch coordinates (rather than space or body frame coordinates; cf. Section 2.3.1) 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. 2.2.

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. 2.2, this is the position of the yellow particle in the circle.

The code also stores the velocities of the particles in two ways, due to the differential shear across the patch: the “wrapped patch velocity” (velocity #1, with shear matching position #1), and the “unwrapped patch velocity” (velocity #2, with shear matching position #2).

Each of these velocities 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. 2.2 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, one 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 to 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 my implementation, I find it is only necessary to store one P_y , corresponding to position and velocity pair #1.

However, once one introduces non-redundant velocities and positions for each particle, one immediately encounters a problem: how should the other types of position and velocity be reconstituted 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.

My 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, the model finds the most massive particle, and

uses its position as a reference point (since that particle is most likely to be closest to the COM—but the method will work regardless). 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. 2.2), 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, the code calculates 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.

I made a critical assumption in order to make this data reconstruction method possible: I 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 showed that obvious numerical artifacts appear in visualizations 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, I must simply make the patch larger.

2.3.3 Aggregate Orientation

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

The Eulerian equations of solid-body motion (Eqs. (2.2) and (2.3)) 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. I 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 my patch coordinate system rotates, sliding out of sync with the inertial space frame, I 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, my aggregates must precess at the orbital frequency of the patch (Ω).

I handle this orientation issue when transforming to and from body coordinates before applying Eqs. (2.2) and (2.3). Consider an integration over the time interval dt . First rotate the aggregate into the space frame by rotating the transformation matrix Λ 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. After the Euler equations are applied, transform from the body frame to the patch, and derotate Λ into the patch frame's new orientation, using an angle $-\Omega(dt + t)$.

In practice, I 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 I rotate into the "space" frame is zero degrees. This saves computational work. Thus the above procedure reduces to simply rotating Λ 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 I 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 I stored the spins in the patch frame, I 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 handling, after a collision is detected).

I recognize the inconsistency between my 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 I discuss required modifications to aggregate acceleration, torque, and stress calculations when in a non-inertial frame. I 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 of Perrine et al. (2011), as described in the following.

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. (1.1)) 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 Perrine et al. 2011).

In contrast, to compute both torques and stresses on rigid aggregates in the rotating frame, I require the accelerations on each particle without Hill's terms—since those add accelerations that do not torque a body directly. I 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, the model invokes Eq. (36) in Per-

rine et al. (2011) to calculate the torque on an aggregate. Note that I 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 properly torque non-spherical aggregates to align their long axes 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).

I 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) in Perrine et al. 2011), 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 I 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](#)).

I 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 con-

siderations for aggregate spin. The code adds 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, the code integrates the aggregate(s) involved forward to the collision time. This involves rotating Λ by the angle $-\Omega dt$ to compensate for frame rotation, integrating the Euler equations for aggregate spin, and drifting the COM. The code then resolves the collision, including any potential merging or fragmentation events. Ideally, it would then trace back the COM position and derotate Λ 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, the code must check for further collisions, and `pkdgrav`’s collision-search algorithm assumes that all particle positions are valid for the start of the step. Thus, in order to detect further collisions, the positions for the aggregate’s constituent particles are back drifted (undoing the frame rotation as well)—but the aggregate’s COM is not. If another collision is detected, the code then integrates the aggregate forward (over the time interval between the newly detected collision and the time the aggregate was last updated), as mentioned above. Again, once the collision is resolved, the code updates the time at which the aggregate’s new information is valid, and the process repeats.

Once all collisions have been resolved, aggregates finish the step by integrating forward (taking an Euler step, drifting the COM position, and rotating Λ) over whatever time

remains in the step. (If there are no collisions for an aggregate during a timestep, it simply takes a full step forward.)

2.3.6 Overlapping Particles and Correction 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, the code extrapolates 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. (2.2) and (2.3) 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 be overlapping. This is an unfortunate inevitability of this strategy—even a quartic expansion of the collision-prediction equations cannot provide the exact collision circumstances (though smaller timesteps help to reduce errors).⁵

Three options to account for particle overlaps are available in the code, which are called 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

⁵An iterative procedure to predict and revise collision times may reduce these errors—a possible future feature of the code.

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 move a particle 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 spuriously.

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. I 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.

I have performed a suite of tests using the repel strategy (the only implemented strategy that avoids numerical artifacts): 70–90% of all overlaps in typical simulations occur 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. I do not feel these overlaps are harmful to the results of my model, as the overlap simply imposes a minor larger-than-average mass density to that region of the aggregate, which mimics a physical variation in material density.

Chapter 3

Cohesion Simulations

3.1 Parameters

Here I describe the parameters used for all 100 simulations presented in this work as part of the cohesive particle project, starting with those that were consistent throughout all runs.

All particles had an internal density of 0.5 g/cm^3 to represent porous ice.

Following Porco et al. (2008), I used the speed-dependent coefficient of restitution (ϵ) law of Borderies et al. (1984), using a v^* of 0.001 cm/s , with zero surface friction. This law has ϵ decrease with increasing impact speed, and this 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 .

The critical opening angle (θ_{crit}) for the gravity tree was 0.5 radian, and my timestep was approximately 5 seconds. I 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 to 50 seconds, I find my timestep is a reasonable compromise between accuracy and computation time. (A timestep of 50 seconds resulted in significant errors in the model's outcome; on the other hand, a timestep of 0.5 seconds yielded fewer overlaps, but otherwise no systematic differences, at the cost of over 10 times greater computation time.)

The dimensions of my patches were constant within the A or B ring simulations: approximately 880 by 350 meters for the A ring, and 695 by 280 meters for the B ring, with the longer axis in the azimuthal direction (the aspect ratio of 2:5 was chosen to accommodate the pitch angle of the gravity wakes). For my fiducial choices of surface mass density (Σ) for each ring (500 and 1000 kg/m² for the A and B ring, respectively), these patch dimensions correspond to 4 by 10 critical wavelengths (λ_{crit} —see Section 1.3).

But changing Σ can change the number of particles (N); for simplicity, and to keep N tractable (as runtime scales—at best—with N^2 , due to geometrically increasing collision frequency), I chose to keep my patch size static when varying Σ (rather than forcing larger patches with increases in λ_{crit}). This means that at higher surface densities (and thus larger λ_{crit} 's) there is a higher risk for wake self-interaction (between the simulated patch and its replicated ghost cells). However, my densest patches are at worst double the Σ of my fiducial—and were thus 2 by 5 λ_{crit} —and my experience is that 2 by 5 λ_{crit} provides reliably similar results to the fiducial patch size, so I feel that this is an acceptable compromise between numerical accuracy and runtime.

I placed my simulated patches at two orbital semimajor axes to mimic B and A ring conditions, at 100,000 and 136,530 km from Saturn, respectively (with orbital periods of 8.96 and 14.3 hours). The majority of my simulations used a 1-m-radius monodisperse particle population. (The choice to use relatively large particles keeps N to a tractable $\sim 75,000$ in the majority of my runs.) For comparison purposes, some simulations were

performed with polydisperse particles, with a size range of 0.8–1.2 m, and using a power-law size distribution exponent $\alpha = -3$, assuming the form $n(R) \propto R^\alpha$. These choices result in nearly uniform N between the polydisperse and the monodisperse cases. (Note that the average particle radius in the polydisperse case is approximately 0.96 m, which means that v_{esc} will be $\sim 4\%$ smaller on average in the polydisperse cases than in the monodisperse.) In either case, a size distribution of aggregates emerges as the simulations proceed, resulting in changes to the overall effective α .

Particles were placed randomly within the patch, with no initial aggregates or spin, 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, I 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.) Each run was carried out to about 9 simulated days (15 orbits for my A ring runs, 24 orbits for the B ring runs). This gives plenty of time for studying the equilibrium state of the system, as equilibrium is established within approximately 5 orbits (as determined when properties such as velocity dispersion and optical depth level off with time).

Most of my simulations completed in 8–11 days on 16 processor cores, depending on the interparticle collision rate—but the highest Σ runs took drastically longer. In fact, my B ring simulation with the highest Σ (2000 kg/m²) has only reached $\sim 70\%$ completion after 93,000 CPU hours, as of the printing of this dissertation. (Its data is nonetheless included in this work—with caveats—as it is important to test the response of the aggregate population at higher values of Σ when possible.)

My simulations explored the parameter space by holding certain values constant while varying a single parameter, in order to isolate the effects of that parameter. Different suites

of simulations explored the effects of the merge and fragmentation limits, bond strength, and mass surface density (see Tables 3.1 through 3.4). Most of these parameters are not experimentally well-constrained, especially for 1 m ice spheres, providing primary motivation for this work. However, when available, I chose best-guess fiducial values and reasonable ranges for my parameters based on observations or theoretical considerations.

My fiducial merge limit was $0.5 v_{esc}$, or 0.27 mm/s, which was based on the results of Hatzes et al. (1991), who found the critical sticking speed for frosty 2.5 cm ice spheres to be ~ 0.3 mm/s. I recognize that there likely is a relationship between merge limit and radius, but I assumed a fiducial merge limit of $0.5 v_{esc}$ for this work; this is necessary, as I am constrained to larger particle sizes (see above) for which no cohesion data exists (to my knowledge).

I have no such experimental reference for the impact fragmentation limit, so I chose two physically reasonable values of 1 and $2 v_{esc}$ (0.53 and 1.06 mm/s). In my model, the fragmentation limit cannot be lower than the merge limit, so to choose a lower fiducial fragmentation limit would confine my possible merge limits during parameter sweeps. Thus the lower limit of $1 v_{esc}$ is perhaps more physically grounded, but the larger limit of $2 v_{esc}$ allows me to explore higher merge limits.

For my fiducial bond strength I reference the only extraterrestrial Solar System ice whose strength has been studied (to my knowledge): comets. I chose 100 Pa, as this seems consistent with the analysis by Sekanina and Yeomans (1985) for the strength of Comet Brooks 2. Asphaug and Benz (1996) found a weaker upper limit (~ 5 Pa) for the strength of Comet Shoemaker-Levy 9, while Richardson et al. (2007) found a higher upper limit of $\sim 10^{3-4}$ for Temple 1. My fiducial appears to be a reasonable midpoint.

The reader should keep in mind the discussion in Section 1.3—namely that `pkdgrav` assumes $\beta = 2$ for strength conversions, and one should divide by the proper β in order to translate into the appropriate real strength for the bonding model being considered.

That is, if β_{real} is small, the real strength (S_{real}) is obtained by plugging the simulated parameter strength ($S_{pkdgrav}$) into the relation $S_{real} = S_{pkdgrav}/\beta_{real}$.

For example, if we assume that bonding is occurring in the rings via the frost layer overlap model provided by Hatzes et al. (1991), we would like to translate my parameter strengths into real strengths appropriate for that model. My approximate β_{real} for that model is 10^{-3} (see Appendix A.1), so a $pkdgrav$ bond strength parameter of 100 Pa translates to 10^5 Pa in the frosty ice layer cohesion model.¹

While my fiducial fragmentation limits and strength seem plausible, I believe that they are also roughly consistent with one another. Appendix A.2 provides a rough order-of-magnitude estimate relating the impact speed between two bodies (v_{imp}) and the pressure exerted by that impact (P_{imp}):

$$\left(\frac{1}{\beta}\right) \left(\frac{R}{1\text{ m}}\right) \left(\frac{\rho}{1\text{ g/cm}^3}\right) \left(\frac{v_{imp}}{1\text{ mm/s}}\right) \left(\frac{1\text{ s}}{\Delta t}\right) = \left(\frac{P_{imp}}{1\text{ Pa}}\right) \quad (3.1)$$

with R the radius of the (equal-size) impacting bodies, ρ the internal mass density of the bodies, Δt the timescale of the collision, and assuming $\beta \ll 1$. (Note that this equation is a duplicate of Eq. (A.8), included here for convenience.) This relationship allows one to set the impact fragmentation limit (a speed) to be consistent with the bond strength (a pressure), and enabling a convenient scaling between these two failure limits.

Assuming $R = 1\text{ m}$, $\rho = 0.5\text{ g/cm}^3$, and $\Delta t = 0.01\text{ s}$ (see Appendix A.2), we find that an impact fragmentation limit of 0.5 mm/s is approximately consistent with an aggregate strength (tensile and shear) of 25 Pa ($\beta = 2$), or 10^5 Pa ($\beta = 10^{-3}$). Thus for $pkdgrav$, where $\beta = 2$, a fragmentation limit of 0.5 mm/s is within an order of magnitude of my fiducial strength (100 Pa).²

¹I acknowledge that this strength is higher than the upper limits found observationally for comets. This value is only a fiducial, and I have conducted a strength-varying suite in order to test the population's sensitivity to strength. See Section 4.3 for a discussion of my results regarding the strength parameter.

²As above, if I translate my parameter strength into a real strength based on the model of Hatzes et

My fiducial Σ values for the A and B ring were 500 kg/m^2 and 1000 kg/m^2 , respectively, producing approximately equal computational load for each location within the rings. My A ring value is based on the consistent results found both by Colwell et al. (2009) and Robbins et al. (2010). The true B ring Σ may be much larger than 1000 kg/m^2 (Robbins et al. 2010), but even doubling to 2000 kg/m^2 turned out to be extraordinarily expensive (as mentioned above, my 2000 kg/m^2 simulation has only reached $\sim 70\%$ completion as of the printing of this thesis). Thus I maintained the B ring fiducial of 1000 kg/m^2 in the interest of completing a large number of simulations to produce parameter sweeps—being sure to also provide a suite studying the aggregates’ state at a range of Σ ’s.

For details regarding the remaining parameters, see Tables 3.1 through 3.4, in which runs are organized into suites with one changing parameter and all others constant (as indicated by dashes). The tables provide an “Index” for each run, which is a shorthand for the parameters of interest in that simulation. The first component of the index is (typically) A or B, indicating A or B ring nominal parameters (Saturnian distance and surface density). (Note that this component is omitted in the final suite in Table 3.4.) The second component indicates the parameter varied in that suite, with M for merge limit, F for fragmentation limit, Str for strength, and Sig for surface density; when the merge limit is varied, I specify the fragmentation limit used (1 or $2 v_{esc}$), as I employ two different fragmentation limits in such runs. The third component is the nature of the innate particle population, either monodisperse (“mono”—all particles having radius of 1 m) or polydisperse (“poly”—radii of 0.8–1.2 m, with $\alpha = -3$). Finally, the fourth component is a convenient run number within the suite. When I reference a suite throughout this work, I simply state the three first components of the indices of the runs in that suite; for al. (1991), my choice for the fiducial strength is somewhat higher than the cometary data suggests—but it remains approximately consistent with my choice for the fragmentation limit.

Index	M (v_{esc})	F (v_{esc})	S (Pa)	Σ (kg/m ²)	a (km)
A-M1-mono-1	0.01	1.0	10 ²	500	136530
A-M1-mono-2	0.05	-	-	-	-
A-M1-mono-3	0.1	-	-	-	-
A-M1-mono-4	0.175	-	-	-	-
A-M1-mono-5	0.25	-	-	-	-
A-M1-mono-6	0.375	-	-	-	-
A-M1-mono-7	0.5	-	-	-	-
A-M1-mono-8	0.75	-	-	-	-
A-M1-mono-9	1.0	-	-	-	-
A-M1-poly-1	0.01	1.0	10 ²	500	136530
A-M1-poly-2	0.05	-	-	-	-
A-M1-poly-3	0.1	-	-	-	-
A-M1-poly-4	0.175	-	-	-	-
A-M1-poly-5	0.25	-	-	-	-
A-M1-poly-6	0.5	-	-	-	-
A-M1-poly-7	1.0	-	-	-	-
A-M2-mono-1	0.05	2.0	10 ²	500	136530
A-M2-mono-2	0.25	-	-	-	-
A-M2-mono-3	0.5	-	-	-	-
A-M2-mono-4	0.75	-	-	-	-
A-M2-mono-5	1	-	-	-	-
A-M2-mono-6	1.25	-	-	-	-
A-M2-mono-7	1.5	-	-	-	-
A-M2-mono-8	1.75	-	-	-	-
A-M2-mono-9	2.0	-	-	-	-

Table 3.1: First in a series of tables of simulation parameters, with each row representing a simulation. The first column is the simulation index, as explained in Section 3.1. Note that the index includes both the code for the ring region under study (either “A” for A ring, or “B” for B ring), and the inherent distribution of indestructible particles (either monodisperse [“mono”] at 1 m, or polydisperse [“poly”] at 0.8–1.2 m, $\alpha = -3$). The next 3 columns in the table are the bonding parameters: merge limit (M), fragmentation limit (F), and strength (S), in either units of v_{esc} (0.53 mm/s, the escape speed from a spherical particle with radius of 1 m and density 0.5 g/cm³), or Pascals. The last column is the mass surface density (Σ). The total combined CPU time for all the runs represented on these four tables is over 540,000 CPU hours.

Index	M (v_{esc})	F (v_{esc})	S (Pa)	Σ (kg/m ²)	a (km)
A-F-mono-1	0.5	0.5	10 ²	500	136530
A-F-mono-2	-	0.75	-	-	-
A-F-mono-3	-	1.0	-	-	-
A-F-mono-4	-	1.25	-	-	-
A-F-mono-5	-	1.5	-	-	-
A-F-mono-6	-	1.75	-	-	-
A-F-mono-7	-	2.0	-	-	-
A-F-mono-8	-	2.5	-	-	-
A-F-mono-9	-	3.0	-	-	-
A-F-mono-10	-	3.5	-	-	-
A-F-mono-11	-	4.0	-	-	-
A-F-poly-1	0.5	0.5	10 ²	500	136530
A-F-poly-2	-	0.75	-	-	-
A-F-poly-3	-	1.0	-	-	-
A-F-poly-4	-	1.25	-	-	-
A-F-poly-5	-	1.5	-	-	-
A-F-poly-6	-	1.75	-	-	-
A-F-poly-7	-	2.0	-	-	-
A-F-poly-8	-	2.5	-	-	-
A-F-poly-9	-	3.0	-	-	-
A-Str-mono-1	0.5	1.0	10 ⁻⁵	500	136530
A-Str-mono-2	-	-	10 ⁻⁴	-	-
A-Str-mono-3	-	-	10 ⁻³	-	-
A-Str-mono-4	-	-	10 ⁻²	-	-
A-Str-mono-5	-	-	10 ⁻¹	-	-
A-Str-mono-6	-	-	10 ⁰	-	-
A-Str-mono-7	-	-	10 ¹	-	-
A-Str-mono-8	-	-	10 ²	-	-
A-Str-mono-9	-	-	Infinite	-	-
A-Sig-mono-1	0.5	1.0	10 ²	125	136530
A-Sig-mono-2	-	-	-	250	-
A-Sig-mono-3	-	-	-	375	-
A-Sig-mono-4	-	-	-	437	-
A-Sig-mono-5	-	-	-	500	-
A-Sig-mono-6	-	-	-	625	-
A-Sig-mono-7	-	-	-	750	-
A-Sig-mono-8	-	-	-	1000	-

Table 3.2: Second in a series of tables of simulation parameters; this table continues to specify the A ring runs. See Table 3.1 for details.

Index	M (v_{esc})	F (v_{esc})	S (Pa)	Σ (kg/m ²)	a (km)
B-M1-mono-1	0.0	1.0	10^2	1000	100000
B-M1-mono-2	0.01	-	-	-	-
B-M1-mono-3	0.05	-	-	-	-
B-M1-mono-4	0.1	-	-	-	-
B-M1-mono-5	0.25	-	-	-	-
B-M1-mono-6	0.5	-	-	-	-
B-M1-mono-7	0.75	-	-	-	-
B-M1-mono-8	1.0	-	-	-	-
B-M2-mono-1	0.05	2.0	10^2	1000	100000
B-M2-mono-2	0.1	-	-	-	-
B-M2-mono-3	0.25	-	-	-	-
B-M2-mono-4	0.5	-	-	-	-
B-M2-mono-5	0.75	-	-	-	-
B-M2-mono-6	1.0	-	-	-	-
B-M2-mono-7	1.5	-	-	-	-
B-M2-mono-8	2.0	-	-	-	-
B-F-mono-1	0.5	0.5	10^2	1000	100000
B-F-mono-2	-	0.75	-	-	-
B-F-mono-3	-	1.0	-	-	-
B-F-mono-4	-	1.25	-	-	-
B-F-mono-5	-	1.5	-	-	-
B-F-mono-6	-	1.75	-	-	-
B-F-mono-7	-	2.0	-	-	-
B-F-mono-8	-	2.5	-	-	-
B-F-mono-9	-	3.0	-	-	-
B-Sig-mono-1	0.5	1.0	10^2	250	100000
B-Sig-mono-2	-	-	-	500	-
B-Sig-mono-3	-	-	-	750	-
B-Sig-mono-4	-	-	-	1000	-
B-Sig-mono-5	-	-	-	1250	-
B-Sig-mono-6	-	-	-	1500	-
B-Sig-mono-7	-	-	-	2000	-

Table 3.3: Third in a series of tables of simulation parameters; this table specifies most of the B ring runs. See Table 3.1 for details. (Note that B-Sig-mono-7 is included here, even though it remains incomplete.)

Index	M (v_{esc})	F (v_{esc})	S (Pa)	Σ (kg/m ²)	a (km)
B-Str-mono-1	0.5	1.0	10^{-5}	1000	100000
B-Str-mono-2	-	-	10^{-4}	-	-
B-Str-mono-3	-	-	10^{-3}	-	-
B-Str-mono-4	-	-	10^{-2}	-	-
B-Str-mono-5	-	-	10^{-1}	-	-
B-Str-mono-6	-	-	10^0	-	-
B-Str-mono-7	-	-	10^1	-	-
B-Str-mono-8	-	-	10^2	-	-
a -mono-1	0.0	0.0	0.0	500	100000
a -mono-2	-	-	-	-	110000
a -mono-3	-	-	-	-	120000
a -mono-4	-	-	-	-	130000
a -mono-5	-	-	-	-	140000

Table 3.4: Fourth in a series of tables of simulation parameters. See Table 3.1 for details. The first suite is the final set of B ring runs, and the second suite consists of runs without cohesion to study the effects of Saturnian distance (a) independent of aggregation.

example, the first suite in Table 3.1 is A-M1-mono (i.e., nominal A ring settings, with variable merge limit, fixing the fragmentation limit to 1 v_{esc} , and monodisperse initial particles).

Note that some parameters are duplicated across different runs; to save computation, I often used a single run in many suites (e.g., entries A-M1-mono-7, A-F-mono-3, and A-Sig-mono-5 are the same simulation). However, I did occasionally take the opportunity to rerun a simulation (with the same parameters and initial conditions) in order to check its results for consistency (e.g., entries B-M2-mono-4 and B-F-mono-7 are independent simulations with the same parameters). This helped me check for uniform results between the two different computer systems I used for my runs.

3.2 Data Analysis

Once the particles equilibrate (as determined by examining dynamical properties such as filling factor at the midplane and velocity dispersion), a simulation proceeds for approximately 10 orbits, during which I sample the state of the system ~ 10 times per orbit—measuring, for example, the percentage of free (unbonded) particles, the number of aggregates (including 2-particle “dumbbells”), the radius of the largest aggregate (R_{max}), the slope of the size distribution (α), the physical optical depth (τ_{phys}), and the particle vertical scale height (h).

I estimate the effective radius (R) of an irregularly shaped aggregate by first finding its principal axis lengths ($a \geq b \geq c$) then computing the radius of the volume-equivalent sphere as $R = \sqrt[3]{abc}$. Once the sizes of the aggregates are known, I bin the data into a log-log histogram, and measure the slope of the line (see Fig. 3.1). I remove from the slope all bins with bodies less than 1.2 m in radius; I found that the results from the monodisperse and polydisperse runs showed improved agreement when I used this method, as it essentially ignores all unbonded bodies in both types of simulations. For example, Fig. 3.1 shows that particles in the first bin are far overrepresented in the power law, and would drastically skew the fit if included. Moreover, in the polydisperse case, there is more than one overrepresented bin, as there are unbonded bodies of many sizes in those plots—thus the polydisperse and monodisperse cases would include *different* biases if these bins were included. So my convention now is to ignore bins that include unbonded particles, or $R \leq 1.2$ m. It is possible that these unbonded particles are overrepresented because of their artificial indestructibility—these particles would perhaps in reality break down into a distribution of ever-smaller particles, which may in equilibrium match with the rest of the aggregate population. But this is beyond the scope of the model at this time.

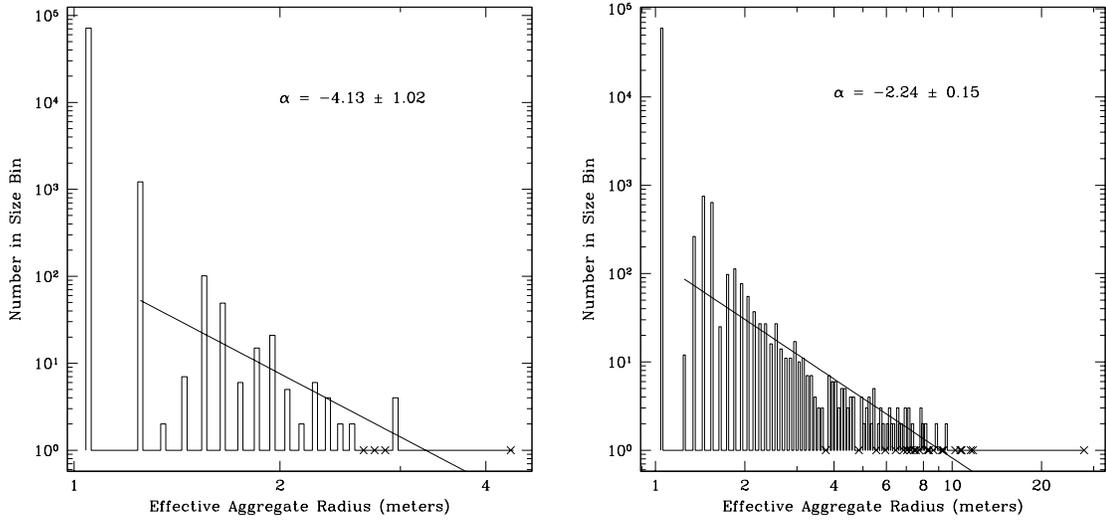


Figure 3.1: Two example incremental size distributions of effective radii of bodies, chosen from snapshots of two simulations. The left pane shows the final state of the aggregate population in run A-M2-mono-1, demonstrating the relatively steep size distribution (and smaller number of aggregates) that comes with weaker cohesion. The right pane shows the final state of the aggregate population in a run with stronger cohesion, A-F-mono-5, in which larger bodies flatten out the size distribution. Bins are 0.1 m wide, X’s indicate bins that contain exactly one body (since logarithmic axes cannot natively distinguish between bins containing zero or exactly one object), and the result of the least-squares fit to the histogram’s slope on this log-log plot is given as α (with $1-\sigma$ uncertainty), providing the power-law exponent in the size distribution function $n(R) \propto R^\alpha$. Note that in each plot there is a gap between the largest body and the next-largest body (and the location where the slope intercepts the x -axis). In this work, when I discuss the largest aggregate (R_{max}), I refer to the single body to the far right of these plots. (If I were to instead define R_{max} as the x -intercept, for example, my results for R_{max} would decrease by a significant amount ($\sim 33\text{--}50\%$), so this is an important distinction.) Also note that the fit does not include bodies of radius less than 1.2 m (see discussion in Section 3.2).

The physical optical depth (τ_{phys}) is determined by firing “test rays” into the particle field at random locations normal to the plane of the ring, and computing the probability that the rays pass through the ring. (Specifically, $e^{-\tau_{phys}} \equiv n/N$, where n is the number of test rays that pass through the ring and N is the total number of rays fired. Thus τ_{phys} is infinite if $n = 0$, and $\tau_{phys} = 0$ if $n = N$.) Finally, h is the RMS height of all of the particles in the patch.

Once each quantity is determined for each sampled timestep, I take the average and

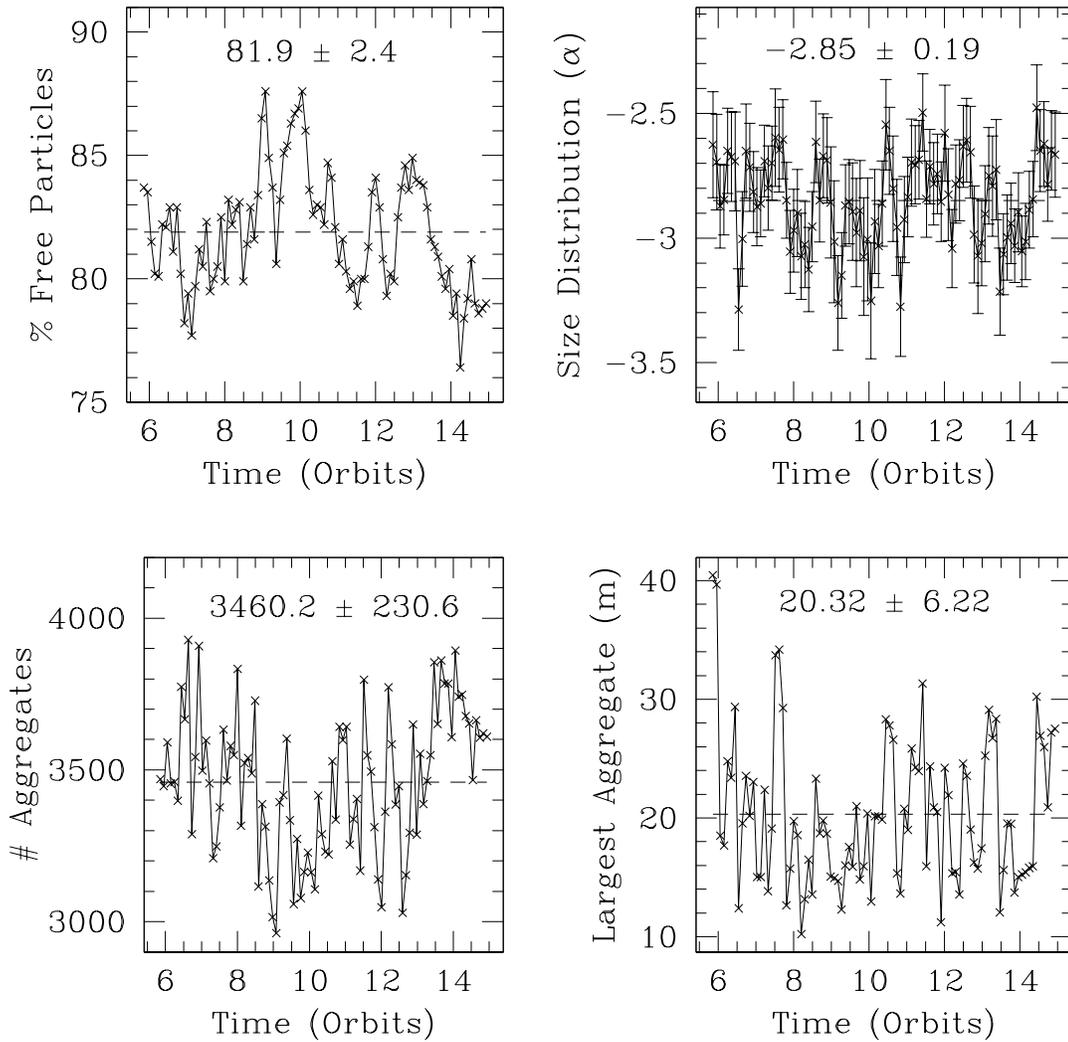


Figure 3.2: Shown here is the time evolution of four properties of the aggregate population in run A-M1-mono-8 (see Table 3.1). Displayed are the final 9 orbits of the 15-orbit simulation, demonstrating the system’s highly variable equilibrium. The upper-left panel is the fraction of unbonded particles 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 upper-right panel shows the slope of the size distribution, α , taken from plots such as Figure 3.1. 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 combines each point’s error with the variance in the time-averaged data to obtain the value displayed on the plot—though most of the quoted error in α results from time-variability.

standard deviation of these quantities over time. This is necessary to remove the natural variability of all of these quantities even when the system is in equilibrium (see Fig. 3.2). 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). Figure 3.2 demonstrates that time-averaging helps to filter out these rapid changes in aggregate evolution. Scatter in τ_{phys} and h can be attributed to wake formation and destruction, and similar scatter in the aggregate properties is caused by stochastic bursts of aggregate formation. These processes are related, as aggregates tend to form readily as wakes collide, creating zones of low encounter speeds and high densities, as discussed in Section 4.1.

Note that this averaging procedure may help my results more closely match observations. While my simulations study a small patch of ring material for a long temporal baseline (approximately a week), observations typically image a large portion of ring material all at once, studying a large spatial baseline at one instant. I believe that these two techniques should yield similar results, as each technique removes natural variability by averaging and estimating error from internal variability in the sample.

Note that my calculation of the size distribution of bodies in my simulations only includes bonded aggregates, and that any loosely bound gravitational aggregates (such as wakes) do not affect my determination of α . This may be in contrast to what observers have measured, as those observations are based on light interacting with the rings in total, which cannot distinguish between gravitational and bonded aggregates as easily. I have not attempted to include gravitational aggregates into my data analysis; this is an area of future work.

Chapter 4

Results and Discussion of Cohesion

Simulations

Snapshots from four selected simulations are presented in Figs. 4.1 and 4.2. The results of all simulations (whose parameters are detailed in Tables 3.1 through 3.4), analyzed as discussed in Section 3.2, are plotted in Figs. 4.3 through 4.10. See the figure captions, and the discussion below, for further details.

4.1 The Lifetime of a Typical Aggregate

Here I provide a brief description of the growth and destruction of a typical bonded aggregate. As discussed in Section 3.2, aggregates form chiefly in the high-density and low-relative-speed environment of the gravitational wakes. Then, as mentioned in Section 1.3, the wakes are short-lived, and dissolve on orbital timescales, due to the shearing nature of the disk. The dissolved wakes leave behind newly created non-spherical aggregates, which then are free to collide (bouncing, merging, or fragmenting), either with other aggregates or free particles. Many survive until coming into contact with another gravitational wake—at which time they are typically destroyed down to their constituent

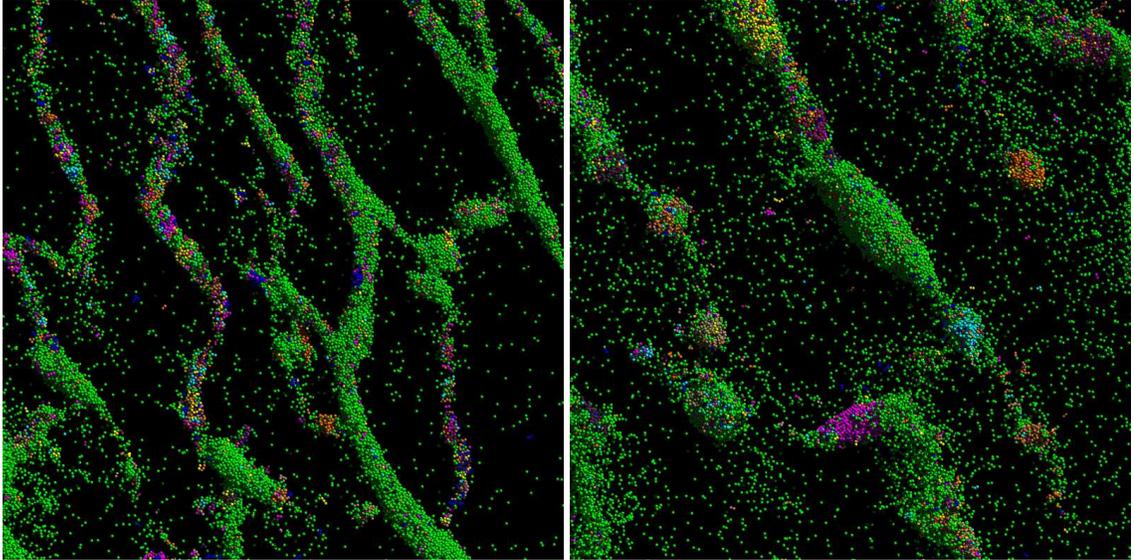


Figure 4.1: Snapshots from A ring runs with moderate cohesion (A-F-mono-4, left; cf. Table 3.2), and high cohesion (A-M2-mono-9, right, cf. Table 3.1). For reference, the merge limit in the left pane is near the Hatzes et al. (1991) value, but the merge limit in the right pane is 4 times “stickier.” Viewing geometry is as in Fig. 1.1. These patches are approximately 350 meters on a side as pictured, so this view only displays approximately half of the simulated region. To enhance aggregate visibility, particles bonded into aggregates (shown in non-green colors) are drawn overtop of non-bonded particles (in green)—thus, aggregates are never hidden behind a cluster of unbonded green particles. Conversely, clusters of green particles are unbonded gravitational aggregates; see for example the large object near-center on the right pane. Note the finer wake structure and abundance of small aggregates in the left pane, but coarser wake structure and larger embedded aggregates in the right pane. Strong bonding evidently not only enables larger aggregates, but begins to disrupt wake structure.

particles, due to the encounter speeds involved. 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}). As wakes form on orbital timescales, the typical aggregate lifetime is an orbital period as well.

If the wakes are very tightly packed together (e.g., Fig. 4.2), the lifetime of an aggregate is even shorter, due to the decreased time between annihilating impacts.

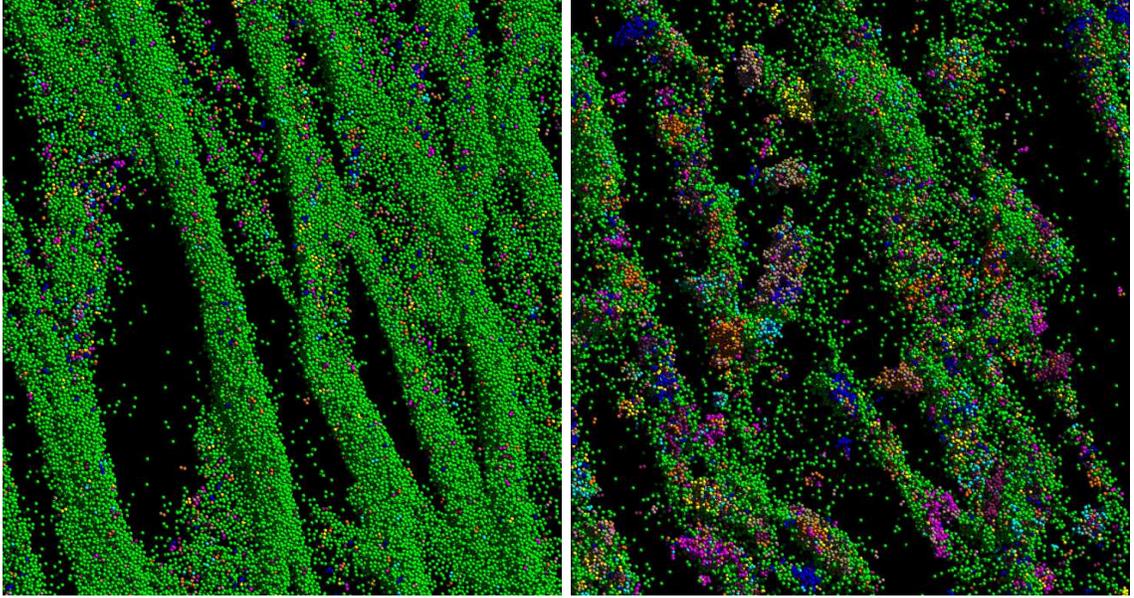


Figure 4.2: Pictured here are two simulations based in the B ring using the same cohesion parameters as in Fig. 4.1, showing moderate cohesion (B-F-mono-4, left) and high cohesion (B-M2-mono-8, right). Viewing geometry and colors are identical to Fig. 4.1. Note the drastic difference in structure between the two panes, in comparison to the two panes in Fig. 4.1—this demonstrates that there may be a relationship between the physical optical depth and cohesion parameters in the B ring, which is not present in the A ring (see Section 4.8).

4.2 Effects of Surface Density

One of the surprising discoveries of the Cassini mission was that the mass surface density (Σ) of the main rings is not as well constrained as previously believed (see Robbins et al. 2010). Consequently, simulations of Saturn’s rings should ideally consider a range of Σ values.

Figure 4.3 shows my parameter sweep over Σ in both the A and B ring regions. The overall conclusion is that changes in Σ do not change my results with respect to matching observations of α and R_{max} with my data—but there are subtle effects. The number of free particles increases with higher Σ , indicating that higher collision rates tend toward liberation of more free particles rather than trapping more particles into aggregates. The

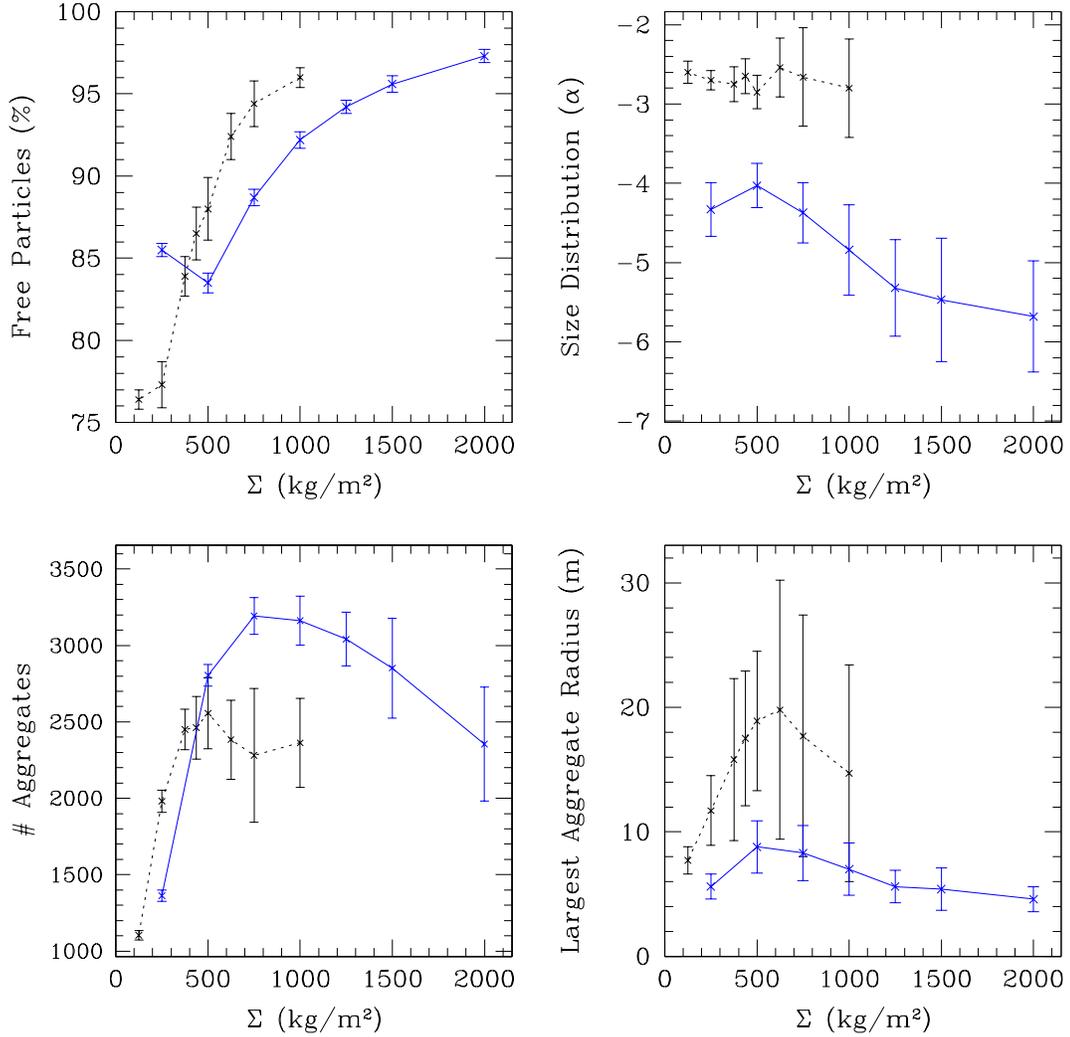


Figure 4.3: Ring mass surface density (Σ) vs. four averaged, equilibrated aggregate population properties calculated as described in Section 3.2. All points represent separate simulations. As in Figure 3.2, the upper-left pane shows the number of unbonded particles in the simulation, as a percentage. The upper-right pane shows the average α , with propagation of error used to combine the errors for each snapshot (from the least-squares fit; see Fig. 3.1) with the variation of α over time. The lower-left pane shows the number of bonded aggregates (including 2-particle “dumbbells”), and the lower-right shows R_{max} . Suite B-Sig-mono is the solid line (blue), and A-Sig-mono is dotted (black); see Tables 3.2 and 3.3. Note that the rightmost points in the solid lines refer to an incomplete simulation (see Section 3.1), and thus carry greater uncertainty than the other points (as they are the result of approximately half of the time-averaged data as the other points).

power law index as measured from the simulations is steady in the A ring, with widening errorbars with larger Σ due to higher impact rates (and thus more stochastic aggregate formation). The average R_{max} peaks in the center of my A ring range, but the errorbars are large enough to assume a constant value of 10–20 m.

However, in the B ring, α appears to become steeper with increases in Σ , moving from approximately -4.5 to -5.5. While α appears to saturate at a value of approximately -5.5 at the highest surface densities, additional simulations at high Σ would be required to discover if α stabilizes or continues to fall at ever larger Σ —unfortunately, these simulations are very computationally expensive. On the other hand, R_{max} appears roughly constant with Σ in the B ring, at around 5–10 m.

Thus my A ring results appear insensitive to Σ , but caution is needed when interpreting my B ring results for α , given the large uncertainty in Σ there.

4.3 Dominant Fragmentation Mechanism: Collisions or Stress?

Figure 4.4 shows the results of my suites studying the response of the aggregate population to variations in the strength parameter, for both the A and B rings. In both regions of the rings, the plots show a sensitivity to the strength below a critical value, and essentially no sensitivity to the strength above that value (to within errorbars). This critical strength, for both the A and B rings, appears to be 10^{-2} Pa (for $\beta = 2$).

Evidently, strength fragmentation is not a significant breakage mechanism above the critical value, as the aggregate populations do not respond to changes in strength above this threshold. For further evidence of this, see the A ring simulation with infinite strength (marked at 10^4 Pa on the figure): the results are similar to every simulation at and above the critical strength. That is, above the critical strength, the bonds are essentially infinitely

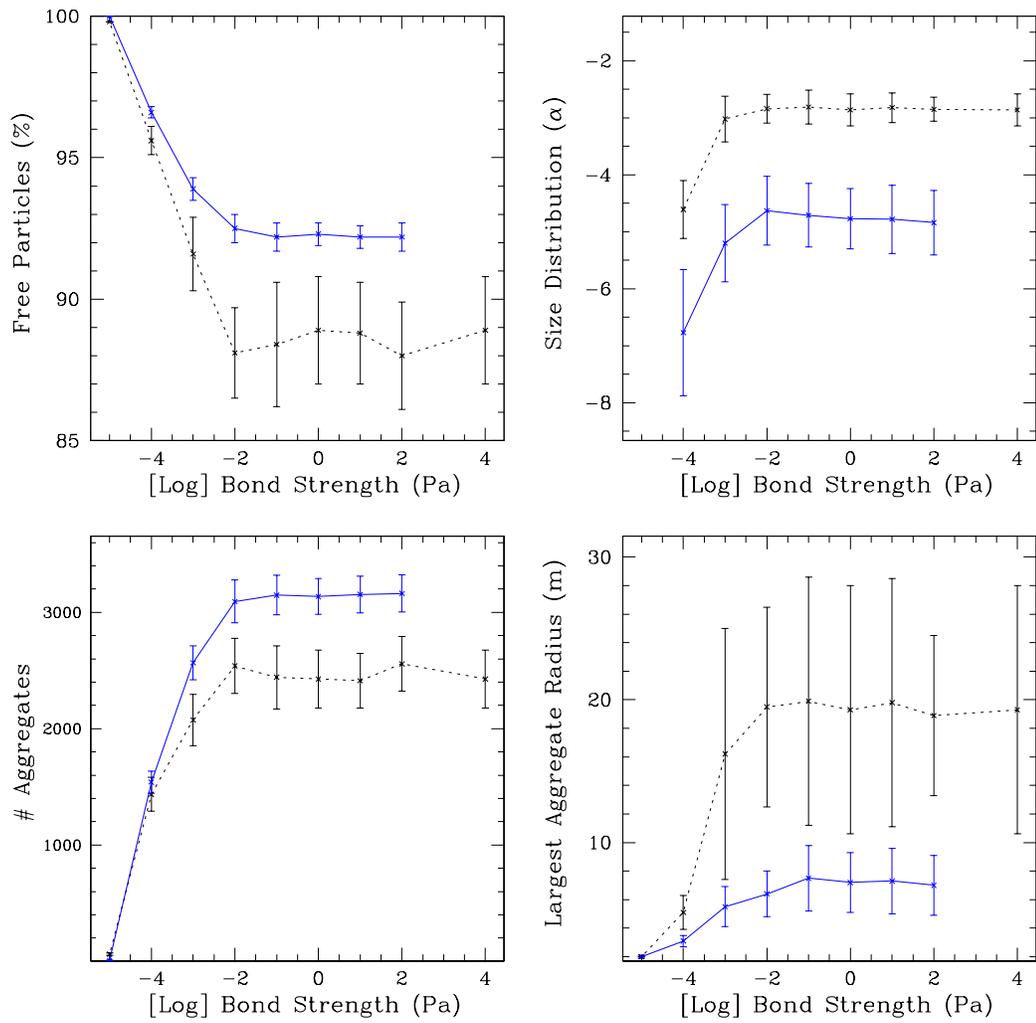


Figure 4.4: Bond strength (essentially an acceleration threshold for fragmentation) vs. the four equilibrated aggregate population properties described in Fig. 4.3. Note that the x -axis is in log space, and that the point plotted at 10^4 Pa represents infinite strength. Lines represent suites B-Str-mono (solid, blue) and A-Str-mono (dotted, black); see Tables 3.2 and 3.3.

strong, as the typical stresses encountered in these regions of the rings are weaker than the bonds.

In the absence of acceleration-stress fragmentation, the only remaining fragmentation mechanism is impact fragmentation. Thus, when the strength is at or higher than the critical strength, the limiting factor for aggregate size is impact fragmentation, which is governed by the fragmentation limit. If below the critical strength, both stress and impact fragmentation can occur.

The majority of my runs set the strength to 100 Pa ($\beta = 2$), as motivated by comet observations (see Section 3.1), so my runs are located in the impact-fragmentation-dominated regime. But these cometary results are only upper limits, and I am aware that the true strength of these comets may be orders of magnitude below my choice of fiducial strength. However, my fiducial strength is 4 orders of magnitude larger than my experimentally determined critical value for the dense rings; thus real ring particles may still lie in the strength regime that is insensitive to strength fragmentation.

By extension, the discovery of the existence of a critical strength implies that stresses like tides and aggregate spin are not very important to the discussion of weak cohesive bonding. If bonding is weak enough to allow Saturn's tides to break bonds (strength under 10^{-2} Pa, for $\beta = 2$), then, according to my relationship between impact pressure and speed (Eq. (3.1)), the impact merge and fragmentation limits would be very low: at least 4 orders of magnitude below the merge and fragmentation limits used in this work. These limits would be so low that essentially every collision would be above the fragmentation limit; thus sticking could not occur often, and impacts would rapidly break any bond that does form. Thus, the most important cohesive bonding parameters for the rings are the merge and fragmentation limits. This conclusion assumes that there is a (linear) relationship between impact pressure and speed (that is, that the derivation of Eq. (A.8) in Appendix A.2 is appropriate for Solar System ice).

4.4 Constraints on A Ring Bonding Parameters

It was established in Section 4.3 that collisions are the dominant fragmentation mechanism for aggregates in my simulations. Here I examine the results from suites that varied merge and fragmentation limits in order to obtain ranges of those parameters that match observations of the A ring.

I presented some of the observations of α and R_{max} for Saturn's A and B rings in Section 1.1, and these results are quite complex. For simplicity, for both the A and B rings, I set the target range for matching the observations of α at -2.75 to -3, and R_{max} at 10–20 m.

I match my chosen criteria for α and R_{max} reasonably well in Fig. 4.5 at a wide range of merge limits (0.053–0.53 mm/s), using the lower of my two fragmentation limits ($1 v_{esc}$). Figure 4.6, on the other hand, uses the higher fragmentation limit ($2 v_{esc}$), and only satisfies the criteria for α and R_{max} at a single point (0.05 mm/s). Clearly, there is a relationship between fragmentation limit and merge limit: if one value is too extreme, as the fragmentation limit appears to be in Fig. 4.6, the other parameter's range that matches the observations becomes tiny. Ideally, the best range in one parameter opens the largest matching range in the other. An exhaustive search over all possible combinations of merge and fragmentation limits is possible, but it would require far more simulations than I have performed in this work, so I will simply estimate the bounds on the region of the merge-limit/fragmentation-limit plane that best match the observations.

In Fig. 4.5, the monodisperse data matches the target α and R_{max} values with merge limits of 0.053–0.53 mm/s, while the polydisperse data is slightly more restrictive, with matching merge limits of 0.09–0.53 mm/s. In Fig. 4.7, the A ring line matches the observations for fragmentation limits within 0.4–0.7 mm/s. This result fits with the discussion above, as the fragmentation limit in Fig. 4.6 is outside of the matching range quoted here,

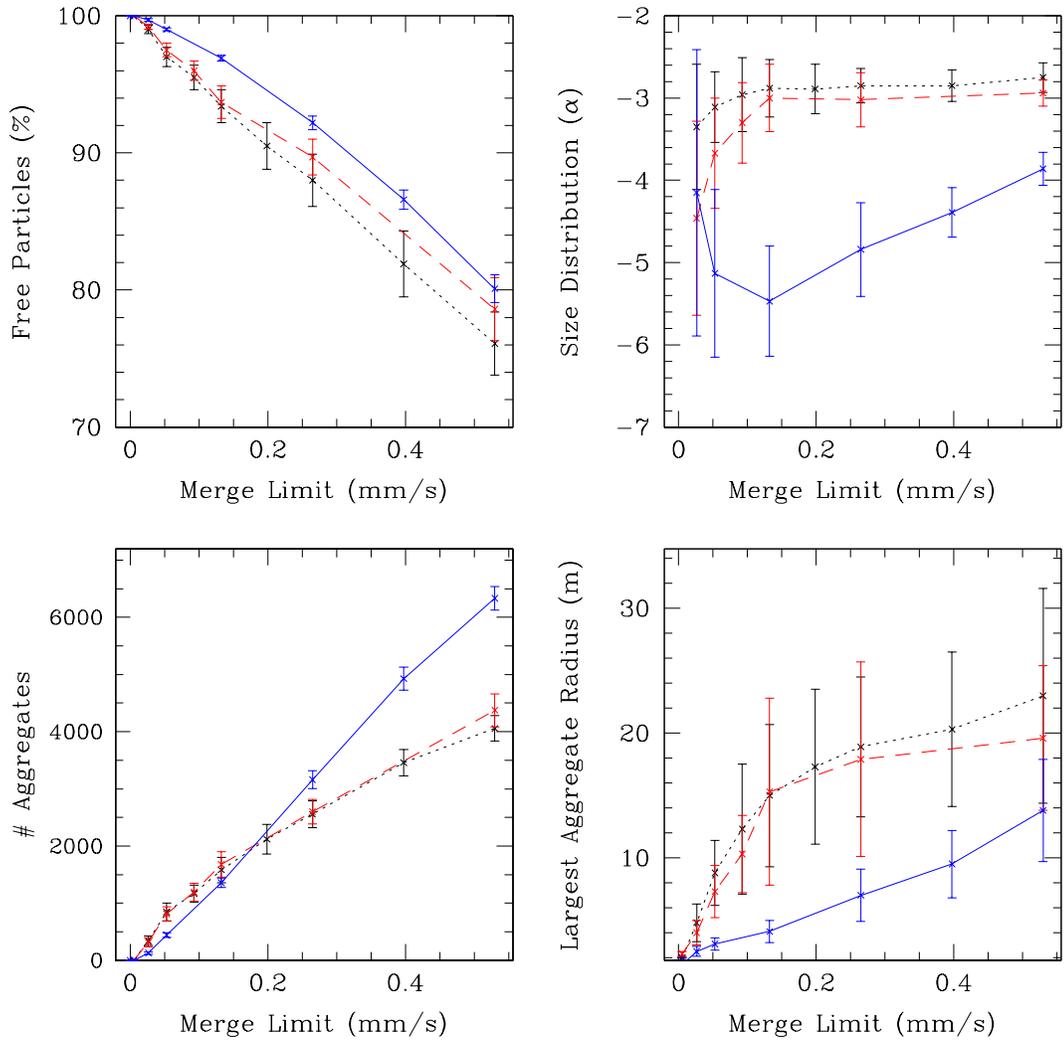


Figure 4.5: Merge limit vs. the four equilibrated aggregate population properties described in Fig. 4.3. All points represent separate simulations with various merge limits, and use the lower of my two fiducial fragmentation limits ($1 v_{esc}$, or 0.53 mm/s). Suite B-M1-mono is the solid line (blue), A-M1-mono is dotted (black), and A-M1-poly is dashed (red); see Tables 3.1 and 3.3. Note that the property with the most time-variability, and the largest errorbars, is R_{max} . This is a symptom of the stochastic formation of very large bodies that are quickly destroyed, resulting in brief spikes in R_{max} vs. time.

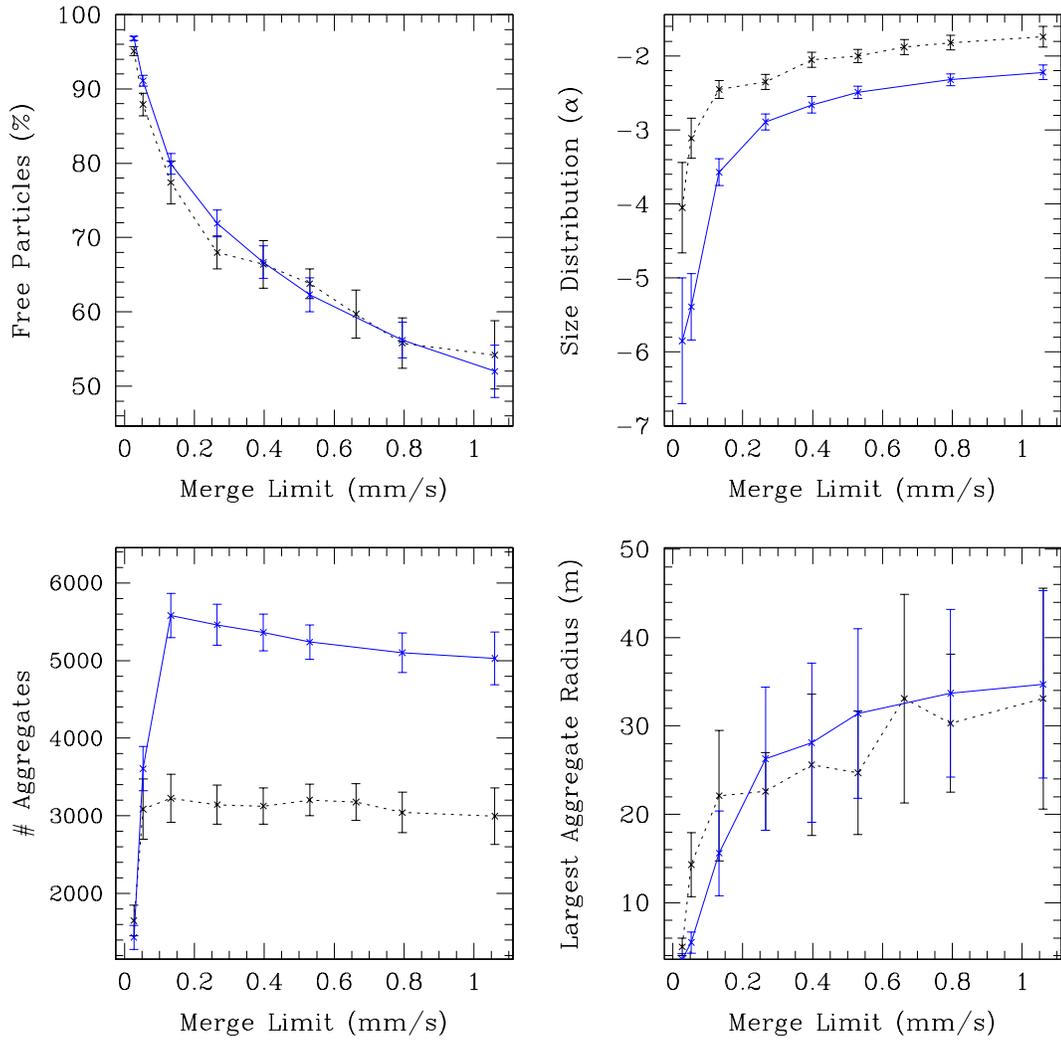


Figure 4.6: Merge limit vs. the four equilibrated aggregate population properties described in Fig. 4.3, but using the higher of my two fiducial fragmentation limits ($2 v_{esc}$, or 1.06 mm/s). Suite B-M2-mono is the solid line (blue), and A-M2-mono is dotted (black); see Tables 3.1 and 3.3.

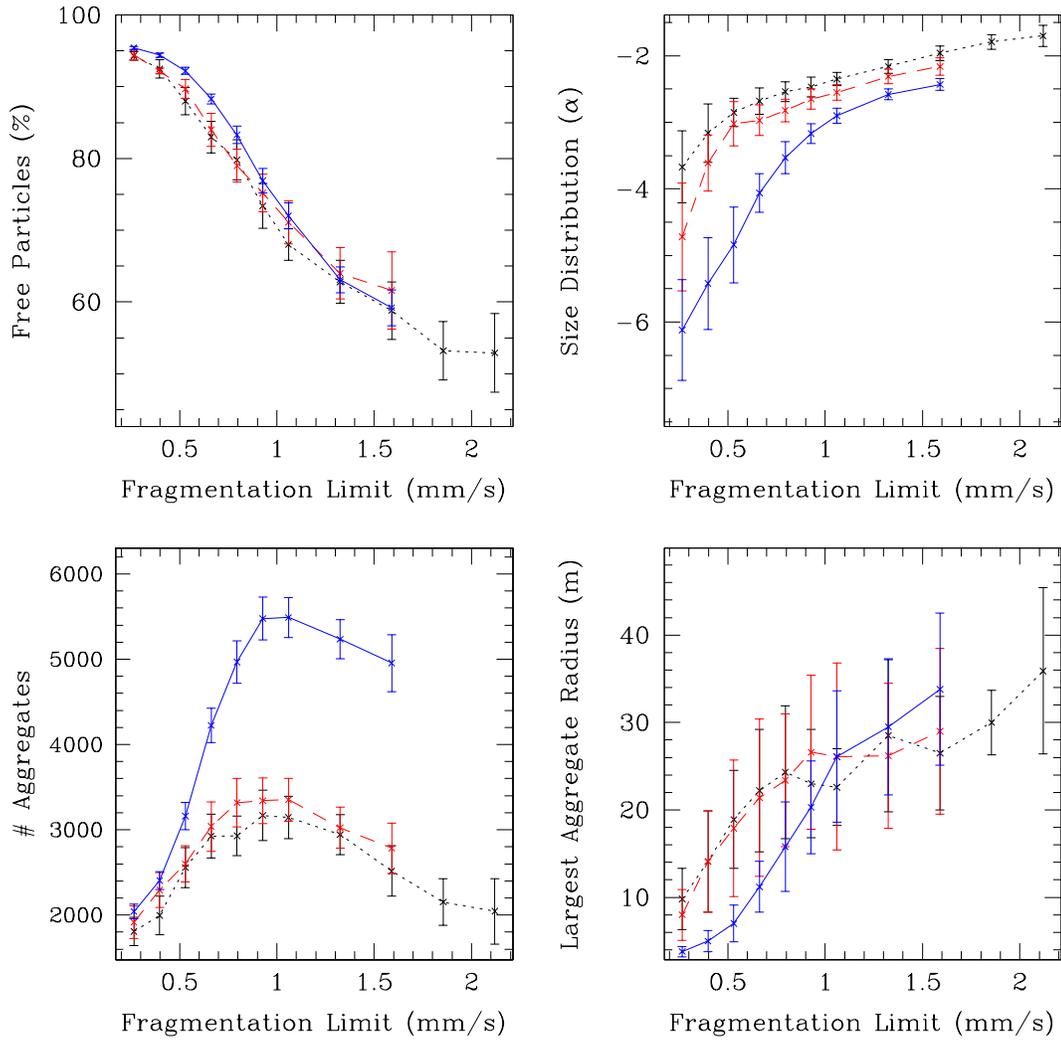


Figure 4.7: Fragmentation limit vs. the four equilibrated aggregate population properties described in Fig. 4.3, with a fixed merge limit of 0.27 mm/s ($0.5 v_{esc}$). Suite B-F-mono is the solid line (blue), A-F-mono is dotted (black), and A-F-poly is dashed (red); see Tables 3.2 and 3.3.

causing a limited range of matching merge limits in suite A-M2-mono.

Thus my estimated A ring range for the merge limit is 0.1–0.5 mm/s, and 0.4–0.7 mm/s for the fragmentation limit. I note that the Hatzes et al. (1991) result of a sticking limit at ~ 0.3 mm/s lies in the center of my range of matching merge limits, lending support to the suggestion that such bonding is possible in the main rings.

4.5 Constraints on B Ring Bonding Parameters

Section 4.4 provided two sets of constraints that I can now apply to the B ring results: a set of observational criteria, and limits on the bonding parameters that matched my A ring results to those observational constraints.

Applying the observational constraints (target $\alpha = -2.75$ to -3 ; $R_{max} = 10$ – 20 m) to the B ring results plotted in Fig. 4.5 turns up no matching parameters. The data in Fig. 4.6 is consistent with the observations at merge limits of 0.25–0.4 mm/s. Lastly, Fig. 4.7 satisfies the target α and R_{max} values at fragmentation limits in the 0.9–1.1 mm/s range. (Evidently, the fragmentation limit of 0.53 mm/s, used for the suite in Fig. 4.5, was too low to satisfy the B ring observational constraints, which is consistent with that limit being below the matching range found in Fig. 4.7.)

Alternatively, if I assume the bonding parameters are constant throughout the rings, and apply the parameter range that matches my criteria in the A ring to the B ring data, I find a curious result: α should be approximately -4 to -5 in the B ring, and R_{max} would be 5–15 m.

What is discovered is that the observational result of consistent α and R_{max} throughout the rings is at odds with the assumption of constant bonding parameters throughout the rings. I discuss possible resolutions to this conflict in Section 4.6, and devote the rest of this Section to discussing why I expect the A and B rings to behave differently in the

presence of constant bonding parameters.

Section 4.3 discussed my belief that stress fragmentation plays a minor role in these simulations when the strength is above a critical value, and thus collisions dominate the state of the aggregate population, so it cannot simply be that, for example, stronger tides in the B ring give rise to different α and R_{max} . Whatever is causing the B and A ring populations to differ must be collisional in nature. The properties of the particles themselves are not changing, so the escape speed (v_{esc}) is the same. It cannot be an effect of Σ , as Fig. 4.3 shows A and B ring simulations with the same Σ with vastly different results. I believe that it is the slight change in the orbital frequency (Ω), and its resulting change in the shear speed (v_{shear}), that is causing such a drastic difference.

At the semimajor axis of my A ring runs, v_{shear} is 0.37 mm/s (for $R = 1$ m particles), which is below their v_{esc} of 0.53 mm/s (for $R = 1$ m, $\rho = 0.5$ g/cm³ particles). But in my B ring runs, Ω is a nearly 60% higher than in the A ring, so v_{shear} becomes 0.58 mm/s, which is now larger than the escape speed. Thus the impact speed distribution shifts to larger values in the B ring than the A, as the greater v_{shear} establishes a higher floor of potential encounter speeds. And if impacts tends to be faster in the B ring, the particles will naturally require higher limits—they will need to be “stickier”—if they are to maintain the same equilibrium sizes and size distributions.

This brings up an important question: since v_{esc} is constant throughout the disk (assuming no drastic changes in particle properties, such as density variations), but v_{shear} falls with distance from Saturn, at what location in the disk does v_{shear} match v_{esc} —that is, where is the turning point outside of which the escape speed is the dominant impact speed? Both speeds are proportional to the size of the body, so it turns out that this location is only dependent on the internal density of the particles. With 0.5 g/cm³, this location is 106,800 km from Saturn’s center,¹ or in the B3 region (Colwell et al. 2009).

¹A $R = 1$ m body with density 0.5 g/cm³ has $v_{esc} = 0.53$ mm/s. I define $v_{shear} = \frac{3}{2}(2R)\Omega$. Solving for Ω

If cohesive bonding is occurring in the rings, and is dominated by collisions as shown in this work, we might observe some changes in the particle population’s properties around this annulus. Given the parameters found for the A ring, I would predict that α is rather steep (about -4 to -5) in the B ring, with somewhat smaller R_{max} (about 5–10 m in radius). As we move out through the rings, we would expect to see a shallower size distribution (less negative α) and growing R_{max} , as the impacts become milder due to a steady drop in shear speed, with a more drastic change around the transition point near the 106,800 km annulus.

It is true that the observations do not seem to agree with my model’s predictions for how α should behave, given constant cohesion parameters, but there are hints that perhaps R_{max} is rising with Saturnian distance. French and Nicholson (2000) show a constant R_{max} with Saturnian distance (20 m), but Voyager observations have hinted otherwise: R_{max} appears to grow from 5 to 10 m in the A ring, and possibly even increases outward through the Cassini division (see Section 1.1). While it is encouraging that the observations for R_{max} from Voyager agree loosely with my findings, I acknowledge that it remains far from a satisfying and firm confirmation of my model’s results.

4.6 Reconciling Observation and Model Results

In Sections 4.4 and 4.5, I found that I match the observational criteria for bonding in the A ring with a merge limit range of 0.1–0.5 mm/s, and a fragmentation limit range of 0.4–0.7 mm/s. The data matches those same criteria in the B ring with a merge limit of 0.25–0.4 mm/s, and fragmentation limit of 0.9–1.1 mm/s. These two sets of limits overlap in the merge limit parameter, but—contrary to my assumption of constant cohesion parameters throughout the disk—differ in their range of fragmentation limits by nearly a factor of

yields 1.8×10^{-4} rad/sec, which is at 106,800 km.

two.

On the other hand, if I were to apply the A ring's cohesion ranges to the B ring, I would expect steeper size distributions there than in the A ring ($\alpha = -4$ to -5), with smaller largest particles ($R_{max} = 5\text{--}10$ m). Here I provide some discussion on ways to reconcile these differences between prediction and observation.

The first explanation to consider is that these results may indicate that impact-limited weak cohesion is not occurring in the rings—that the observed distribution of particle sizes is simply a result of long-term collisional grinding, with particles never growing and coalescing, but merely fragmenting and shrinking over the age of the Solar System. After all, it is a primary goal of this work to explore whether cohesion is compatible with observations of the rings, and this conflict may show that it is not. However, I have managed to satisfy the observational criteria with reasonable parameters, and may even have begun to explain the observed growth of R_{max} with distance from Saturn. Given these promising findings, I will continue discussing ways to resolve the disagreement.

Next, I consider that the bonding parameters can vary in the main rings, and assume that the size distribution and maximum aggregate size are constant throughout the disk. This implies that the particles in the B ring are “stickier,” with a much higher impact fragmentation limit than the particles in the A ring. Given the differences in surface density and orbital frequency between the A and B rings, it is possible that the surface properties of the ring particles would be different, and any frost layers on their surfaces would be potentially compacted and refreshed differently due to the higher collision frequency in the B ring. It may be that the bonds formed between cohering particles in the B ring are twice as strong as those in the A ring, but it seems unlikely that these B ring bonds would be exactly as strong as they need to be to match the equilibrium properties of their counterparts in the A ring. Something fundamental would have to link the strength and fragmentation limits of the bonds to the orbital frequency of the particles.

On the other hand, perhaps the cohesion parameters are constant throughout the disk, and α does change with distance from Saturn. This implies that the observations of the B ring measured different structures than my analysis did—and this seems plausible. Figures 4.1 and 4.2 show that the topography of the A and B rings are very different; perhaps the same wavelength of light, interacting with these very different patches of particles, is sensitive to different structures (P. Nicholson, personal communication, 2010). Perhaps more of the wake structure is measured in the B ring than in the A, where individual bodies are more separated. This potential ambiguity in observational measurements is in contrast to the relative simplicity of my measurements within the model; I have the luxury of counting every body separately and measuring its size—regardless of how deeply embedded in a wake it may be. Thus, until we can analyze the real and simulated patches in the same manner, it will be difficult to know exactly how different our results truly are.

I also point out the possibility that a combination of model parameters may exist that provides a constant α with distance from Saturn—under constant cohesion parameters—that my chosen parameter space did not cover. For example, my range of Σ 's in the B ring may not truly describe the environment there (e.g., Robbins et al. 2010), so my results may not be accurately modeling the true B ring. This seems unlikely however, given my discussion of how changes in v_{shear} are driving these impact-dominated adjustments in the equilibrium properties of the aggregate population (see Section 4.5). I find it unlikely that simply increasing Σ , or using a different coefficient of restitution, could override the inherent changes to v_{shear} that naturally affect the impact speed distribution throughout the disk.

Finally, I must acknowledge that my model remains a first-order attempt at simulating cohesion in planetary rings, and that perhaps my code is too simple, and is missing key physics. While I do not believe my results would change much were I to use a more sophisticated model, the possibility exists.

In summary, my suggestions to reconcile observational results with my model include: cohesion might not be active in the rings, and the observed size distribution arises via another process; particle cohesion parameters vary throughout the rings, which tells us something about the varying surface properties of the particles; cohesion parameters are constant throughout the rings, indicating that the ring observations and my data analysis differ in what they are sensitive to; my parameter space was insufficient to fully study this topic; and my model may not include key physics that could result in better agreement with observations.

4.7 Maximum Aggregate Size in Equilibrium

In this section I discuss the largest average aggregate radii (R_{max}) in my equilibrated ring patches, and address how these findings may apply to “propellers” in the A ring, and the suggestion by Tremaine (2003) that large bonded aggregates may exist in the B ring.

I have observed that it is the impact fragmentation limit that appears to determine R_{max} in a simulation. First, I discussed in Section 4.3 that in my simulations, at my chosen parameters, aggregate fragmentation is dominated by collisional fragmentation, with stress fragmentation playing a minor role in the state of the final aggregate equilibrium population. Then in Figs. 4.5 and 4.6, as each portray suites that use constant fragmentation limits, each suite shows evidence for saturation in R_{max} at higher merge limits. On the other hand, Fig. 4.7 shows consistent growth in R_{max} , with no evident saturation. Thus a rise in fragmentation limit allows R_{max} to rise as well.

As a consequence, I do not expect weak cohesion among ring particles to result in runaway growth: larger aggregates induce faster impacts, as they have higher escape and shear speeds (each of which are proportional to the size of the body—cf. Section 1.3), and so growth will be self-limiting—so long as the fragmentation limit of a body does not

increase with the size of a body (or another mechanism does not come into play at larger sizes to strengthen the body). However, it does not seem likely that icy bodies in the rings should grow stronger with size, as larger (terrestrial) rocky bodies tend to be weaker than smaller bodies (e.g., Housen and Holsapple 1999). If this result applies to ice, the results for R_{max} from my simulations would in fact be too high, as my model makes the simple assumption that the fragmentation limit of a body is constant with size.

The largest aggregates created in my parameter space are approximately 40 m in radius, for both the A and B rings, and only form in simulations that employ the largest fragmentation limits in my parameter space. It is likely that larger aggregates could be made if I extended my parameter range to higher fragmentation limits—but ever higher fragmentation limits lead to stronger disagreement with observations of α , so I do not feel it would be illuminating to provide those models.

Cohesion provides a potential means of constructing large, embedded objects in the dense rings; this is relevant to the discovery of propeller-shaped features in Saturn’s outer A ring, as these features appear to be the perturbed wakes of large (yet unresolved) bodies embedded in the rings. The propellers often appear in multiple images over time (up to 5 years; Tiscareno et al. 2010b), implying that these perturbations are not simply transient fluctuations in the ring material, and appear to be confined to three distinct annuli in the A ring (Tiscareno et al. 2008). These bodies are typically 20–250 m in radius, but could be up to 1–2 km (Tiscareno et al. 2010b). It is unknown whether these bodies are the remnants of a parent body breakup into three well-defined annuli, or whether some mechanism has shepherded these bodies into this configuration. My work explores the alternative, that these propeller bodies are in fact built “bottom-up” out of ring material, rather than “top-down,” as the remnants of a larger body.

My work implies that we can rule out that these bodies are weakly coherent aggregates for three reasons. First, by matching the plausible range of fragmentation limits to

observations of α in the A ring, I have confined the range of possible R_{max} values as well, to approximately 15–40 m. In order to create 250 m (let alone 1–2 km) bodies, I would need much larger fragmentation limits, which would result in α values less negative than -2. Second, these propeller bodies must be long-lived in order to sustain such perturbations over orbital times, and the aggregates in my simulations are created and destroyed on orbital timescales.² Third, my model provides no reason why only three distinct annuli of the A ring would produce propeller bodies. On the contrary, if my model is correct, I would expect to see propellers throughout the A ring, with increasingly larger propellers at larger orbital radii, since R_{max} should increase (according to my model) in that sense (assuming constant bonding parameters throughout the rings). It is possible that there is some threshold size that must be reached before a propeller is created, yet my model provides no reason for these objects to congregate into distinct regions, rather than distributing themselves evenly throughout the ring.

Turning our attention to the B ring: Tremaine (2003) suggests that certain large-scale regions of alternating high and low optical depth in the B ring, ~ 100 km in radial extent, are the result of “shear-free” assemblies of particles—that is, regions of the ring that seem to orbit as a solid body. The present work provides evidence against the suggestion that weak cohesion could build these large structures out of the ring material; while I was able to produce a mix of aggregates, my aggregates’ maximum size is limited by the fragmentation limits permitted by the observational constraints, as discussed above. Tremaine (2003) suggests a plausible strength limit (10^4 Pa) which—were the ice that strong—would allow assemblies of such a size to exist. However, my work shows that the impact fragmentation process is a much more important consideration in regards to

²While it remains possible that larger aggregates may have longer lifetimes due to self-shielding effects, I argue that increases in velocity dispersion from gravitational stirring, and increased shear speed, would destroy these massive bodies just as quickly as the smaller aggregates.

the size and lifetime of these objects, and that due to the steady increases in v_{shear} and v_{esc} that occur as these bodies grow, coupled with the scale of v_{shear} in the B ring, impacts should prevent these objects from growing beyond the scale of tens of meters in the B ring.

4.8 Other Observational Diagnostics for Constraining Aggregation Parameters

To this point, I have focused on constraining bonding parameters to match existing observations of the A and B rings of Saturn (specifically, α and R_{max}). I will now discuss the effects of cohesion parameters on two other potentially observable properties of the rings: the physical optical depth (τ_{phys}) and the scale height (h). If I can show these properties vary with “stickiness,” perhaps I can do more than simply constrain the plausible parameter range.

Figure 4.8 shows the equilibrated τ_{phys} for each of the simulations in Tables 3.1 through 3.3. I first point out that the bottom pane of Fig. 4.8 shows nearly constant τ_{phys} for most Σ values in the A and B ring.³ This agrees with the findings of Robbins et al. (2010), who discussed the nonlinear relationship between optical depth and Σ by providing simulations that showed that one cannot assume that rises in Σ will produce a corresponding linear rise in τ_{phys} (especially at large τ). Thus the following discussion should hold no matter the value of the local surface density (cf. Section 4.2).

It appears from these results that τ_{phys} in the A ring is not sensitive to changes in any cohesion parameter—which unfortunately means that τ_{phys} may not be a useful diagnostic

³I recognize that the last point in pane (h) of Fig. 4.8 does not follow the trends in the data; this is an incomplete simulation (see Section 3.1), and perhaps this point would match the behavior of a slowly rising τ with Σ if it were run to completion.

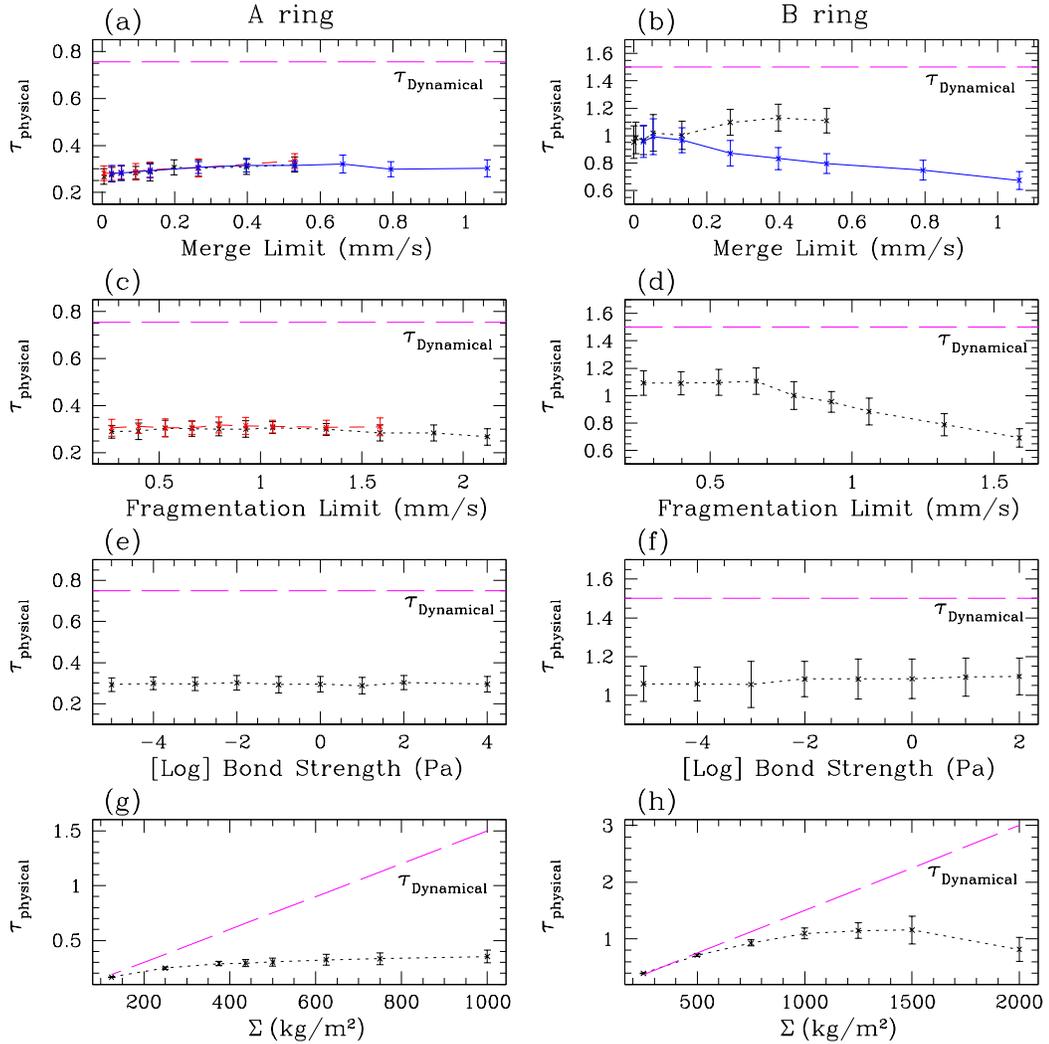


Figure 4.8: These eight plots display the equilibrated physical optical depth in all of the simulations in Tables 3.1 through 3.3 using solid (blue), dashed (red) and dotted (black) lines. Each pane also displays the dynamical optical depth (total cross-sectional area of all particles divided by the patch area) for that simulation for comparison, in a long-dashed (magenta) line. Panes on the left are A ring suites (Tables 3.1 and 3.2), and those on the right are B ring suites (Table 3.3). Pane (a) shows suites A-M1-mono in black, A-M2-mono in blue, and A-M1-poly in red; (b) B-M1-mono in black and B-M2-mono in blue; (c) A-F-mono in black and A-F-poly in red; (d) B-F-mono in black; (e) A-Str-mono in black; (f) B-Str-mono in black; (g) A-Sig-mono in black; and finally (h) B-Sig-mono in black. See Section 4.8 for a discussion. Note that the rightmost point in pane (h) refers to an incomplete simulation (see Section 3.1).

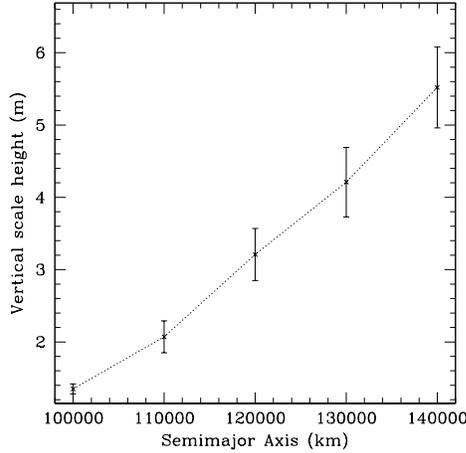


Figure 4.9: Equilibrated vertical scale height (h) from five simulations without cohesion at various orbital distances (a) from Saturn (see Table 3.4). The range in a brackets the two locations used by my A ring (136,530 km) and B ring (100,000 km) simulations, and clearly shows a roughly linear (within errorbars) relationship between patch semimajor axis and scale height.

for cohesion in the A ring. However, in the B ring, τ_{phys} can vary with cohesion parameter by as much as 30%, from $\tau_{phys} \sim 1$ with no cohesion, to $\tau_{phys} \sim 0.7$ at very high merge or fragmentation limits. So, for example, if cohesion parameters were to vary greatly in the B ring, we might be able to observe a change in τ_{phys} .

But why do the same parameters cause changes to τ_{phys} in the B ring, but no changes in the A ring? The answer appears to lie in the scale height of those two regions of the rings. Tests I have performed (see Fig. 4.9) revealed a linear relationship between patch semi-major axis and scale height, even in the absence of any cohesion. (This appears to be consistent with Colwell et al. 2007, who noted that B ring wakes were “flatter” than A ring wakes.) Geometrically, if the material in the A ring is more vertically distributed, there is a greater chance for particles to move above or below one another, reducing the physical optical depth. Thus, even at the same surface densities, an A ring simulation ought to have a lower equilibrium τ_{phys} than an equivalent B ring simulation (which is what I find).

The top pane of Fig. 4.10 shows that h is consistently higher in the A ring simulations,

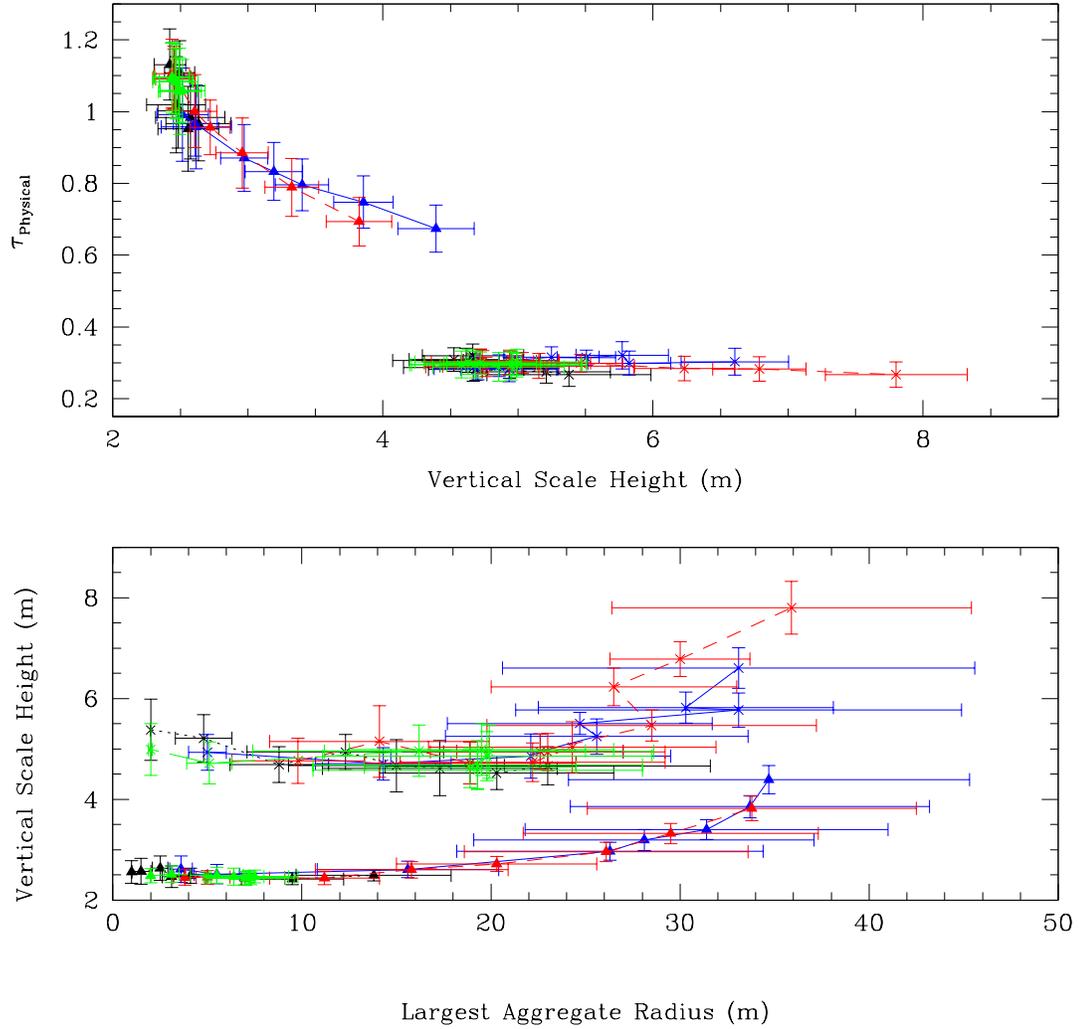


Figure 4.10: Composite plot displaying equilibrated vertical scale height (h) vs. physical optical depth (τ_{phys}) in the top pane, and the radius of the largest aggregate (R_{max}) vs. h in the bottom pane, for many of my simulation suites. Symbols distinguish A ring from B ring suites, with triangles for B ring, and X's for A ring. Line types (and colors) distinguish which parameter was varied: merge limit suites with a lower fragmentation limit ($1 v_{esc}$) are dotted (black), merge limit suites with a higher fragmentation limit ($2 v_{esc}$) are solid (blue), fragmentation limit suites are dashed (red), and bond strength suites are long-dashed (green). See Section 4.8 for a discussion.

while τ_{phys} is always lower in the A ring. In fact, h appears to be so high in the A ring that changes to h do not affect the optical depth (this is presumably because h remains large enough to allow particles to consistently move above or below one another). Conversely, points generally in the upper-left corner of the top pane (B ring data, with weaker cohesion) have low h , and thus such a compact configuration of particles ($h \sim 2\text{--}3$ m) that any increase in h (with stronger cohesion parameters) has a significant effect on τ_{phys} in the B ring.

Why does h increase with stronger cohesion? In the bottom pane of Fig. 4.10, we see that h is excited by increases in the size of the largest aggregate (R_{max}). The scale height appears nearly constant⁴ with changes in R_{max} until $\sim 20\text{--}25$ m, when increases in R_{max} cause h to grow at the highest cohesion parameters (see Figs. 4.5 through 4.7 to map the cohesion parameters to R_{max}).

Thus when cohesion is strong enough to produce aggregates larger than $\sim 20\text{--}25$ m, the particle population responds to the gravitational excitations of these bodies and becomes vertically excited, causing h to rise. If the disk is already thick, as in the A ring, the optical depth does not fall; but if the rings are ordinarily compact, then this rise in h will cause a decrease in τ_{phys} .

My observational criteria included the assumption that aggregates do not grow larger than ~ 20 m in the B ring, and I find that τ_{phys} would only change if aggregates could grow to this size or larger. Thus searching for variations in τ in order to constrain potential cohesion parameters would also potentially shed light on the size of the largest aggregate in the rings.⁵

⁴I note that in the A ring suites, h decreases slightly at lower values of R_{max} until ~ 8 m. I speculate that perhaps cohesion is providing a damping effect (via collisions with zero coefficient of restitution) at these weakly cohesive parameters—an effect that is overtaken at higher cohesion values when the gravitational influence of the large aggregates re-excites the disk.

⁵It would be more direct to observe local variations in h to constrain R_{max} ; however such a study would

4.9 Validity of Monodisperse Starting Conditions

To support the analysis above, I confirm that a major simplification in the model is acceptable—namely, that my use of a monodisperse population of particles (with a uniform 1 m radius) matches the results of my polydisperse comparison cases (radii 0.8–1.2 m, $\alpha = -3$). I ran two comparison suites with polydisperse particles, A-M1-poly and A-F-poly, and their results are plotted with A-M1-mono and A-F-mono in Figs. 4.5 and 4.7, respectively.

To first order, the results are identical, as the points often lie within each other's errorbars. But closer inspection shows that where differences are visible, the polydisperse results are systematically offset from the monodisperse data, with the polydisperse runs having fewer aggregated particles, smaller largest aggregates (R_{max}), steeper size distributions, and more (smaller) aggregates overall. I believe that the presence of the more massive indestructible particles in the polydisperse cases (up to 1.2 m) is increasing the velocity dispersion in those runs, leading to subtle differences between the cases: if the polydisperse particles are excited to larger velocity dispersions, this would result in higher-speed impacts, leading to fewer merging events, and more fragmentation events. This would naturally cause all of the systematic effects observed (though I would not have necessarily expected more aggregates to form, rather than fewer).

I note that this is likely an artificial difference: if these 1.2 m particles were able to fragment, they would not likely have the same enduring effect on the velocity dispersion. This could be tested using a polydisperse population with a smaller maximum

require sub-meter resolution on the vertical structure of the main rings, which is far beyond the capabilities of past or present space missions.

size—however, using smaller particles causes a rapid rise in the number of particles (for a simulated region of a fixed size), and I have not attempted such runs at this time.

Thus I assert that these polydisperse comparison runs show that it is valid (within my model) to use monodisperse indestructible particles, and that my results would not differ significantly were I to use a polydisperse initial size distribution throughout my models.

4.10 Comments on the Range of Bonding Parameters

The chosen range of cohesion model parameters (merge and fragmentation limits, and bond strength) appears to cover the most significant and interesting regions of the parameter space that affect the equilibrium aggregate population. In many of the panes in Figs. 4.4 through 4.7, the data show signs of saturation.

The data approach 0% bonding at low strengths and merge and fragmentation limits. At higher merge and fragmentation limits, there is evidence for saturation in the percentage of bonded particles, as well as α and R_{max} . This is almost certainly the result of my chosen parameter range bracketing the common impact speeds of the rings (the particles' escape speeds and the shear speed of the regions of the rings I have modeled). I predict that my figures would look quite different if the impact speed distribution were centered around other values—that is, if the velocity dispersions of the particles were different. So any alternative choice of parameters that could affect the equilibrium velocity dispersion (e.g., the coefficient of restitution) might in turn affect the equilibrium aggregate population.

I also note that my bond strength range is large enough to encompass both the critical strength value (see Section 4.3) and infinite strength (in one A ring simulation).

I believe that I have sampled the most significant portions of the parameter space that are relevant to the A and B rings of Saturn; however, I note that cohesion should

produce different results in other systems with vastly different velocity dispersions, such as protoplanetary disks in the process of forming planets—requiring a different range of parameters not likely covered by this work.

Chapter 5

Other Rings Work

While working toward the formulation of a suitable thesis topic, and completing the resulting project, I've spent time working on other tasks. Section 5.1 gives a summary of work accomplished prior to starting the dynamics of cohesive ring particles project (which forms the bulk of this thesis), and Sections 5.2 and 5.3 briefly discuss side projects that resulted in published papers, on each of which I appeared as second author.

5.1 Review of Work in Ring-Moonlet Interactions

Prior to starting the cohesive aggregation project, my work centered around modeling ring-moonlet interactions—specifically the gravitational (and, to a lesser extent, collisional) influence of a relatively large (~ 4 km-radius) embedded moonlet on the surrounding ring material. This study was based on the Daphnis/Keeler gap system (see Fig. 5.1), and so the work was focused on moonlets large enough to open up a full gap (as opposed to, for example, the smaller “propeller” perturbers, which are not large enough to disrupt an entire annulus of ring material).

This work involved performing local simulations with a massive moonlet located inside a planetary ring gap. To my knowledge, this had not been done before in a full

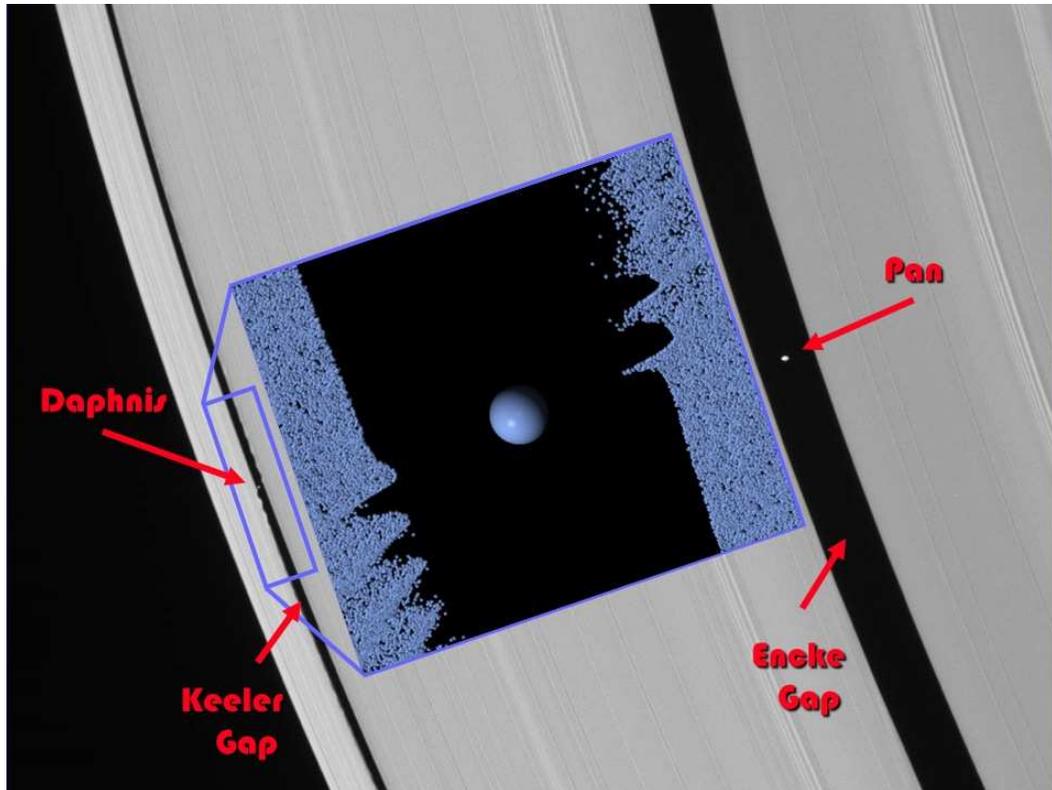


Figure 5.1: A composite view of both simulated and real images of Saturn’s outer A ring. In the background, a Cassini image of the outer A ring, showing the Encke and Keeler gaps, with their embedded moons, Pan and Daphnis. For scale, the Keeler gap is about 40 km wide. The edge waves on the downstream edges of the Keeler gap from Daphnis’ perturbations are clearly visible. Placed overtop is a snapshot from a simulation with parameters similar to the A ring, with a single 3.9 km-radius particle taking the place of the moonlet Daphnis. The view is compressed to enhance the edgewaves produced in the simulation, as indicated by the different aspect ratios between the snapshot and the box surrounding the corresponding location in the Cassini image.

N-body simulation. I soon discovered the challenge this posed, as by introducing a large gravitational perturbation into the local simulation breaks the assumption of generality in the patch model—the patch becomes a unique place in the rings (the place in the annulus in which the moonlet resides), rather than a location that could stand for any neighboring patch. However, if the gravitational disturbance does not extend beyond the simulated boundary, then the model is potentially valid. A simple means of accomplishing this is to use a very large patch; but when larger perturbations are studied, this can require

prohibitively large patches.

Section 5.1.1 discusses the two major alterations that I made to the standard local simulation technique that I utilized in order to perform this work, including changes to the azimuthal boundary conditions (using what I call the “random wrap” technique), and how I chose to model the moonlet within the simulation. Section 5.1.2 then very briefly discusses the theories of edgewaves and gap-clearing that I compared the results of my model against, and presents some results that came out of that work.

5.1.1 Model Alterations

Two means to represent the large perturbing body (the moonlet) within the patch were implemented. The first simply places a large particle (that is approximately 1000 times larger in size than any other particle) in the center of the patch, to represent the moonlet. This strategy has the advantage of full self-consistency within the code, as the moonlet responds to the ring material’s gravity and collisional momentum transfer automatically. One serious complication is that as this particle is treated exactly like all other particles in the patch model, it is replicated into all neighboring patches, potentially complicating the gravitational potential.

To provide a contrasting model, I implemented a capability within the code to represent the perturbing moonlet as an external perturber—a source of gravity that can move independently of the patch, and so can act as a moonlet orbiting about Saturn, and gravitationally influence the particles in the patch externally. The advantages of this strategy include the freedom to place the perturber anywhere (as this potential can be at any location inside or outside of the simulated patch—thus for example representing a very large moon on a resonant orbit), and uniqueness (as the perturber is not artificially replicated by the patch model). The disadvantages include the fact that there is no back-reaction from the particles onto this external potential, as it moves on a fixed path (violating Newton’s

3rd law, though the extent of these errors depends on the mass ratio between the particles and the moonlet), and collisions onto this potential are more difficult to handle (as there is no physical particle for other particles to collide with). Thus, for example, it is not possible to model accretion onto the moonlet.

Each technique can provide a time-variable potential—that is, each is capable of having the moonlet on an eccentric orbit.¹

These two contrasting models allowed me to compare results and determine, for example, that replication of the moonlet, when represented as a particle, did not cause any problems so long as I used a reasonably large patch. This built my confidence in this simple strategy, and so I used it in most of my subsequent work.²

But having a large gravitational perturber in the center of a local simulation creates problems: in order to allow for reasonable patch sizes while ensuring that the embedded moonlet’s perturbation does not unphysically distort the patch material, I created the “random wrap” routine, which erases these perturbations as particles cross over the azimuthal boundary. This ensures that the perturbation is contained to the downstream flow of particles, as the moonlet’s gravitational disturbances are erased as the particles cross from the downstream side of the patch to the upstream side.

This routine modifies the standard periodic boundary conditions: rather than directly reintroducing particles on the opposite side of the patch when they exit, the new procedure “randomizes” particle positions and velocities to erase any dynamical signature of having encountered the moonlet. This attempts to mimic the natural damping effect that a body

¹As it happens, Daphnis was discovered to have a tiny but non-zero eccentricity (Jacobson et al. 2008), and I was able to reproduce the Cassini images with my model (see Fig. 5.6).

²Using the large particle to represent the moonlet was especially useful when I studied accretion rates onto the moonlet, and granular dynamics of that accreted material as it flows onto the moonlet’s surface. It turns out that this material accumulates near the moonlet’s Lagrange points, and then bursts off the surface in sudden and massive flows of particles, somewhat similar to landslides.

experiences after encountering the moonlet in the real ring system: by the time a particle encounters the moonlet again, it has had sufficient particle-particle interactions to return to equilibrium. My procedure artificially removes the moonlet's perturbations without having to perform a computationally expensive simulation over hundreds of thousands of kilometers of ring material.

However, the physical accuracy of this technique requires that we know, in advance, the equilibrium conditions of the ring particles. The following assumptions were made in order to produce an "equilibrated" populations of randomized particles. First, I assumed that the equilibrated (upstream) particles' velocity dispersion matches that of an equilibrated region of the rings that is free of any perturbations from a moonlet. Thus, when a particle wraps azimuthally, the method assigns it a random velocity chosen from a velocity dispersion that was measured in a separate equilibrated and unperturbed simulation. (Note that randomizing a particle's velocity, as chosen from an equilibrated velocity dispersion, is essentially the same as replacing the particle's eccentricity with a tiny random value.)

The randomized particle must have a new position as well, as the torque from the moonlet's perturbation likely changed the particle's semimajor axis. At first, I had assumed that the gap would always remain a fixed width, and thus the radial positions of randomized particles would lie within fixed boundaries given by the user. After observing some peculiar behaviors, I realized that imposing an unphysically fixed gap width in the patch was a poor assumption, as the width of the gap should scale in some way with the mass of the moonlet, and this imposed boundary condition was holding the system out of equilibrium.

In order to model a self-consistent gap width, I revised the method with a more sophisticated assumption: when a particle wraps, its new radial position should not be determined randomly between fixed positions, but it should be located at the guiding center

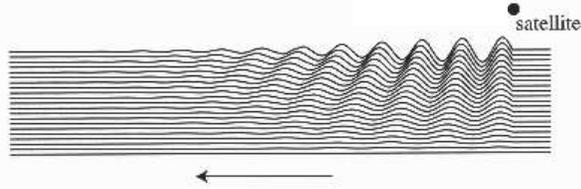


Figure 5.2: Schematic diagram of the linear theory of edgewaves. As material overtakes the moonlet (via disk shear), streamlines become eccentric, which manifests as a wave that dampens as the streams interact and dissipate energy. In this linear theory, the wavelength of these edgewaves is $\lambda = 3\pi\Delta a$, with Δa the distance from the gap edge to the moonlet. Figure reproduced from Murray and Dermott (1999).

of the perturbed particle that exited the patch. In this manner, most of the perturbation is erased, but the particle flow can respond to the torques of the moonlet in a self-consistent manner. (Specifically, any torque imposed from the moonlet that results in changes to the particle’s orbital semimajor axis is preserved across the boundary, but its eccentricity is still dampened and randomized.) This revised method allowed me to proceed with my work to compare linear theory of gap-moon interactions (see Section 5.1.2).

5.1.2 Early Results: Edgewaves and Gap-clearing

The first goal of the project was to form the edge waves that are common in these moonlet/ring gap systems. Figure 5.2 shows a textbook diagram showing the theory of how these waves are created. The wavelength of these structures is predicted from linear theory to be:

$$\lambda = 3\pi\Delta a \tag{5.1}$$

with Δa the distance from the gap edge to the moonlet. I was ultimately successful in producing these waves (see Fig. 5.3), and will discuss specific results below.

The width of a gap that is opened by an embedded moonlet is determined when the torques from the moonlet, which act on the ring material to force the particles away,

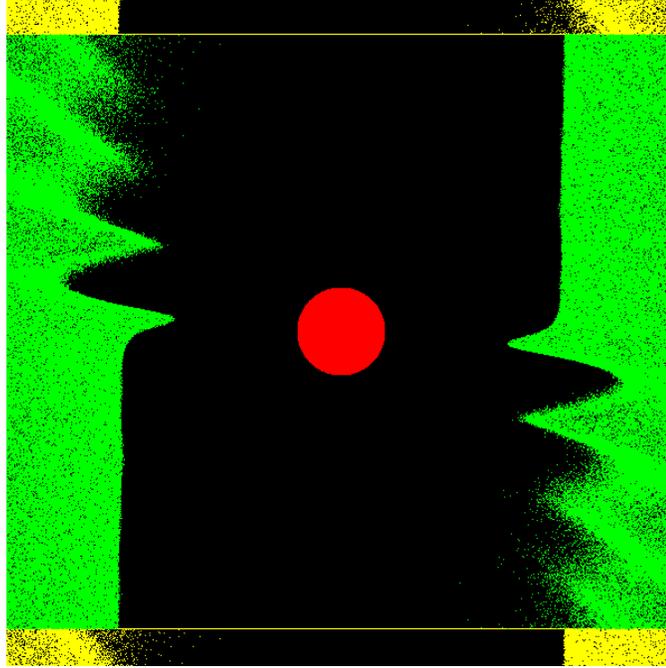


Figure 5.3: A snapshot of a simulation with a 3.9 km-radius moolet placed in the center of a 40 km-wide gap in a patch, producing edge waves along each downstream side of the patch (cf. Fig. 5.2). The orientation is the same as in Fig. 1.1. As thousands of green particles (each $1000\times$ smaller than the red moonlet) shear past the perturber, their orbits become eccentric, and they form streamlines—which then manifest as waves. The patch is 1600 km long, and 120 km wide; the view is compressed vertically to enhance the waves’ visibility. Yellow particles are the replicated bodies that provide boundary conditions. Note that the perturbation from the moonlet is erased across the boundary by invoking the simple random wrap procedure (see Section 5.1.1), leading to a static gap width.

balance the natural viscous spread of the rings to fill in the open gap. Once these two influences are in equilibrium, a stable gap is formed.

Goldreich and Tremaine (1982) developed an analytical expression for the width of a ring gap based on this principle. An adapted version of that expression (appearing in Porco et al. 2007) is:

$$\mu^2 = \left(\frac{3\pi\nu}{0.84\Omega} \right) \left(\frac{\Delta a^3}{a^5} \right) \quad (5.2)$$

with μ as the ratio of the moonlet’s mass to Saturn’s mass, ν the ring viscosity, Ω the moonlet’s orbital frequency, Δa the ring gap half-width, and a the moonlet’s orbital semi-major axis. Plugging in values appropriate for the Keeler gap, leaving only the viscosity

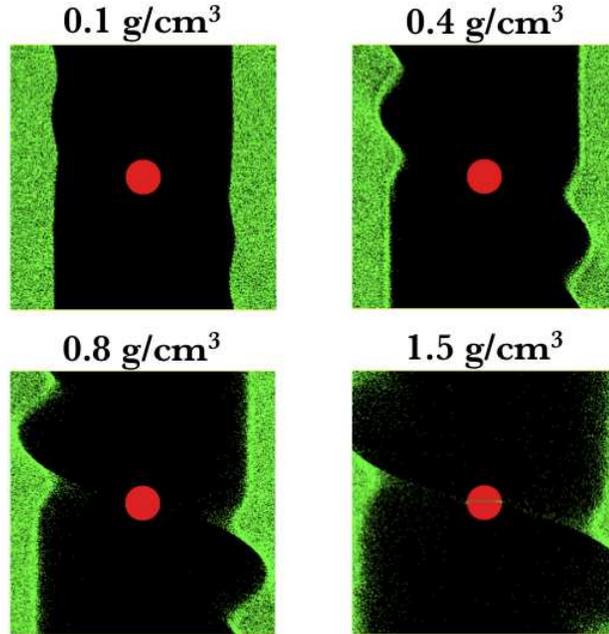


Figure 5.4: Four snapshots of simulations similar to Fig. 5.3, but with the more sophisticated “random wrap” boundary condition technique that places wrapped particles at the guiding center of the exiting particle (see Section 5.1.1). Shown here are four simulations with four different moonlet densities (as labeled), which merely serves to scale their masses, as each moonlet is 3.9 km in radius. The refined boundary conditions allow for the ring material to respond self-consistently to the perturbations of the moonlet. As a result, the variation in mass has a clear effect on the width of the resulting gap, as well as the length of the edgewaves.

and the moonlet’s density as variables, and converting to convenient units, I find the width of the gap as a function of the ring’s viscosity and the moonlet’s density:

$$\Delta a = 99.52(\nu)^{-1/3}(\rho_{moon})^{2/3} \quad (5.3)$$

with Δa again the distance from the gap edge to the moonlet (in km), ν the viscosity (in cm^2/s), and ρ_{moon} the internal density of the moonlet (in g/cm^3).

Using the more sophisticated “random wrap” technique, I observed the correct behavior in models with varying moonlet masses, with more massive moonlets creating larger equilibrium gaps (see Fig. 5.4). I discovered that gap growth is quite rapid, implying that the process of torquing the rings (which causes the ring particles to migrate away from the moonlet) works on a much more rapid timescale than the opposing viscosity (which

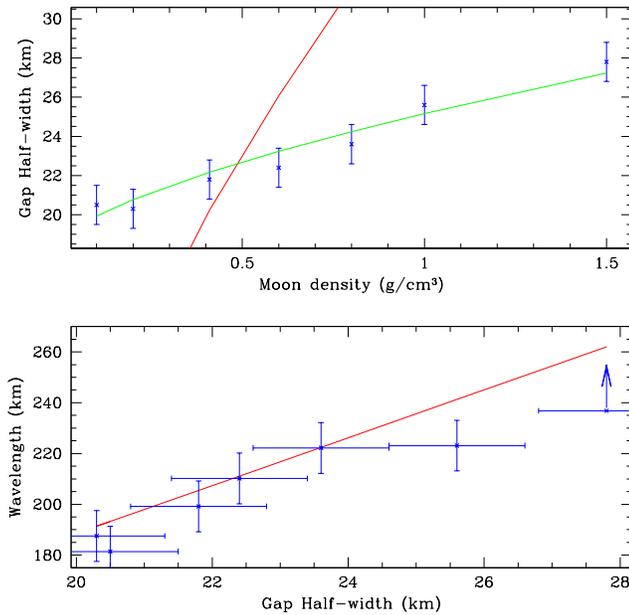


Figure 5.5: Data from seven simulations like those appearing in Fig. 5.4 using the guiding-center-based “random wrap” boundary condition technique (see Section 5.1.1). Blue points are the measured data, showing the gap half-width (distance from the center of the moonlet to the gap edge), the edgewise wavelength (peak-to-peak distance), and the moonlet’s internal density. In the top pane, the red line is the analytical prediction for the gap width (Eq. (5.3)), assuming the viscosity of $20 \text{ cm}^2/\text{s}$ from Porco et al. (2007). The green line is a fit to the data; however the parameters of the fit do not match well to the prediction from Porco et al. (2007) (see Section 5.1.2). The red line in the bottom pane is the prediction from linearized theory (Eq. (5.1)). The relation fits fairly well until the larger perturbations, where it appears the linear theory may break down.

works to keep the gap closed). These results were presented at the 2007 DPS meeting (Perrine and Richardson 2007). Unfortunately, this discovery also implies that the time to reach a stable equilibrium between ring viscosity and moonlet torque is far beyond the scope of my simulations, so I suspect that I never fully reached equilibrium in my models.

Once my simulations appeared to show the proper qualitative behaviors (as both gap width and edgewise wavelength appeared to grow with moonlet mass), I attempted to match my results with Eqs. (5.1) and (5.3) quantitatively (see Fig. 5.5).

The wavelength of the edgewise waves were measured by visual inspection of the distance between wave crests. Often the wave would crest multiple times on the downstream side

of the patch, and so multiple wavelengths could be measured and averaged, improving accuracy. (The simulation with the most massive moon, and thus the largest perturbation, only showed one wave crest, and thus only a lower limit to the wavelength for that simulation is shown in Fig. 5.5.) The gap width was measured visually as well,³ and the errorbars for both the wavelength and gap width were estimated by uncertainty in these on-screen measurements (e.g., taking multiple measurements with the same tool and estimating uncertainty by convolving the variability of the collected data with the uncertainty of the tool itself).

The lower pane of Fig. 5.5 shows fairly good agreement between the prediction of Eq. (5.1) (in red) and the data points (in blue) at lower perturbations (smaller gap widths, from a lower-mass moonlet), but the prediction rises higher than the data at larger perturbations. The data suggests that the linear prediction breaks down at higher moonlet masses, which is reasonable, given that the theory is based on test particles responding to a small perturbation. (On the other hand, given the discussion above regarding the timescale for gap formation, it's likely that these larger perturbations have not yet reached equilibrium. However, equilibrium would likely widen the gap as well, so the data points may not fit to the prediction even given sufficient time to equilibrate.)

The top pane of Fig. 5.5 shows the results comparing variations in moonlet mass (shown here via density variations) to the resulting gap width. The red line shows the prediction from Eq. (5.3), using $\nu = 20 \text{ cm}^2/\text{s}$ (Porco et al. 2007), which does not match my data at all. The green line is a preliminary fit I performed (by hand) to the data using Eq. (5.3), and the resulting fit parameters indicate $\nu = 3340 \text{ cm}^2/\text{s}$, but require an artificial

³However, I later created an automated technique to detect the gap edge, which could be used to refine the data in Fig. 5.5, and may provide a starting point for investigation into this topic by future students. This method first determines the equilibrium particle density in the unperturbed portion of the strip of ring material and then finds the location where the particle density drops to one-half of the average value; this location is labeled as the “edge” of the gap.

offset (y-intercept) of +18.5 km. These highly unusual parameters imply that Eq. (5.2) is not an appropriate model for my early results. At the time of this study, I attributed this disagreement to imprecise measurements of the gap width (which became the motivation to create the automated technique for detecting gap widths mentioned above in footnote 3), and an absence of an equilibrated gap width at the end of the simulations (as I estimated that it would take a computationally impractical amount of time to reach a final equilibrium).

However, this project was dropped in favor of the aggregation project, so this issue was never resolved. These preliminary results are included here so that future students can carry on with the project.

Finally, observations of Daphnis have revealed that this tiny moonlet is actually on an eccentric orbit (Jacobson et al. 2008). By utilizing my model's capability to include a moonlet with non-zero eccentricity (see Section 5.1.1), I was able to reproduce the Cassini images that suggest that Daphnis is on a non-circular orbit (see Fig. 5.6) and recover what was expected: as the moonlet moves on its eccentric path, the gravitational potential acting on the gap edges is time-variable, causing the amplitude of the edgewaves to rise and fall on the same frequency as the moonlet's orbital frequency. By refining the moonlet's orbital parameters and mass, I was able to match well to the images of the Daphnis/Keeler gap system (even predicting Daphnis' orbital phase), and my orbital parameters for the moonlet (modeled eccentricity = 0.000035) agree with those estimated from the observations (observed eccentricity = 0.000033; Jacobson et al. 2008). (These results were presented in Perrine and Richardson 2007.)

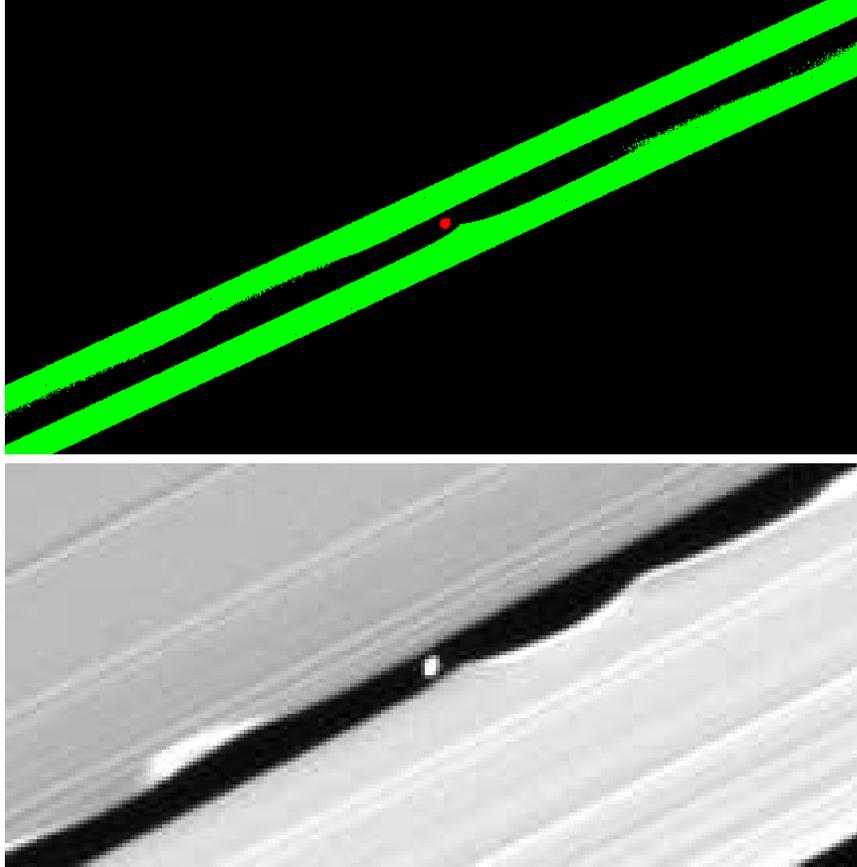


Figure 5.6: A comparison between simulation and observation demonstrating evidence for an eccentric moonlet. The top frame is an example snapshot from a simulation of the Keeler gap/Daphnis system, with the moonlet in the center (red) orbiting with an eccentricity of 0.000035, and the ring material in bands on either side (green). Note that for computational expediency, the simulation only includes a narrow band of ring material on either side of the gap; periodic boundary conditions still apply, but are not shown. The bottom frame is a Cassini image of Daphnis in Saturn's Keeler gap, which has been observed to have an eccentricity of 0.000033 (Jacobson et al. 2008). The viewing angle of the simulation data is adjusted to match the spacecraft's position at the time of observation. Note that the small wave to the lower-right of the moonlet in each image is not mirrored by a symmetric wave to the upper-left. This indicates a time-variable potential; thus the moonlet is on an eccentric orbit.

5.2 An Analytical Parameterization of Wake/Interwake Optical Depth Contrast

Concurrent with the cohesion project, I became involved in another topic in ring dynamics. I provided ring particle simulation data (without cohesion) to Dr. Tiscareno, to which he applied his new analytical technique to parameterize the optical characteristics of a patch of dense ring material (in which gravity wakes reside; see Section 1.3) using only six numbers: 3 optical depths, and 3 weighting factors.

Specifically, the method calculates local densities (and converts these densities into local dynamical optical depths— τ_{dyn}) for every location within a fine grid laid atop an equilibrated patch of particles. (For details, including methods for quality control and validation of data points, see Tiscareno et al. 2010a.) After calculating the density for each grid point, the method creates a histogram showing the frequency with which each density value occurs. Upon applying this method to the provided simulation data, the analysis revealed not a bimodal distribution of optical depths (as had been previously assumed to be the case, corresponding to one optical depth for the wake-centered regions and one for the interwake regions), but a trimodal distribution (with one value of τ_{dyn} for the wakes, and two separate τ_{dyn} values describing the interwake regions). The method then fits three Gaussians to this trimodal data: the sizes of the three fits become the area-based weights, and the central locations of these fits become the three τ values.

By using these six numbers, one can reconstruct the observed τ by finding the area-weighted average of these three optical depths. Further, this work revealed that the optical depth in the interwake regions is not a single value, describing a single photometrically active population of particles, but is an area-weighted average of two distinct populations: one photometrically active and the other inactive. This is significant as most of the area is dominated by this photometrically inactive material.

Though the practical effects of this discovery on the typical analysis of ring observations appears to be minimal (with discrepancies only arising at low opening-angles⁴), this work did shed light on one of the mysteries of the “propeller” objects: it was not yet clear why these objects are seen to be “relative-bright” on both the lit and unlit sides of the rings. This work supports the hypothesis that the presence of the embedded bodies (which perturb the rings to form the propeller features) also locally disrupt the wake structure. This allows material—which would otherwise be locked into dense wakes—to be released into the interwake regions, becoming much more photometrically active, and thus creating bright regions that surround the perturbing objects. As the interwake regions were previously thought to be a single population of photometrically active material, this narrative did not explain why the disruption of the wakes would brighten the rings so dramatically. This work showed that much of the interwake regions are in fact filled with photometrically inactive material, allowing for a much greater contrast in brightness when the wakes are disrupted.

The paper concludes with a discussion of future work to study this process with detailed numerical simulations—simulations that present a significant computational challenge, due to their simultaneous need for small particles (to resolve the gravity wakes) and kilometers of ring material (to encompass the propeller feature and allow for sufficient re-equilibration of the perturbed material). As such, these simulations require an immense number of particles. I hope that these future studies will be enabled by my “random wrap” procedure, which should allow for simulations with notably smaller patches of ring material (and thus fewer particles). Since this method erases perturbations instantly, the simulated region can include only those features directly under examination—as even a factor of two reduction in patch size could be critical to the feasibility of such a study.

⁴The opening angle is the viewer’s inclination; it is zero when viewing the rings edge-on, and 90 degrees when face-on.

5.3 A Symplectic Method for Local Simulations

Hill’s equations of motion for particles in a local simulation (Eq. (1.1)) require modification in order to remain symplectic with a leapfrog integrator (due to the presence of the particle’s velocity—a first derivative—in the equations of motion). Without such changes, orbital elements, like eccentricity, tend to drift unphysically—increasing spuriously over time.

The improved integrator involves splitting the Hamiltonian into exactly integrable parts—in this case, into “free-particle” and “momentum-dependent” terms. It is in the “momentum-dependent” term that the new canonical quantity, P_y , which is similar to a particle’s angular momentum, is introduced. (P_y is ordinarily conserved in the absence of external influences, but must be updated following any collisions or other impulsive events—see Section 2.1.)

Professor Thomas Quinn (U Washington) developed these improvements to the equations of motion, and informed Prof. Richardson and myself of the revised equations. I then incorporated the modifications into my patch-based aggregation model within `pkdgrav` and tested them, thereby helping verify the instructions in Quinn et al. (2010) for how to implement the technique into a complex code. See Section 2.1 for specific considerations regarding the integration of this method into `pkdgrav`.

Chapter 6

Summary and Conclusions

6.1 Summary

I 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.

I have detailed how I modified the existing `pkdgrav` model’s rigid body capabilities to function in a local, rotating frame. I also discussed the frequency and severity of particle

overlaps in the model, and the three implemented resolution strategies, only one of which (the “repel” method) I have found to be adequate for the physical requirements of the systems modeled here.

This 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 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; cf. Section 4.3)—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.

I then used this new code to explore the possibility that cohesion, via the sticking process studied by Hatzes et al. (1991), might be occurring in the A and B rings of Saturn. I presented 100 simulations that explore the parameter space that includes the three cohesion parameters (merge limit, fragmentation limit, and bond strength), the surface density, and the orbital distance, comparing both polydisperse and monodisperse initial particles populations. See the following section for the conclusions of this study.

Finally, I discussed a sample of other projects undertaken prior to and concurrently with the cohesive aggregates project. My early rings work was to model ring-moonlet interactions—specifically the Daphnis/Keeler gap system in Saturn’s outer A ring. I outlined the new modeling techniques used to enable this work (i.e., the “random wrap”

boundary condition), and discussed my results in matching observations and linear analytical theory to the simulations, as they pertained to edgewaves and gap-clearing. I was able to verify that—at lower perturbations—the linearized relationship between edgewave wavelength and the half-width of a gap in which a moonlet is embedded works well, but shows signs of breaking down at larger perturbations. I was unable to verify an analytical relationship between the width of that gap and the mass of the moon, however.

I also briefly discussed my role in two side projects (which resulted in published work). I provided Dr. Tiscareno with a suite of simulations that he analyzed with a new technique that specifically measures the optical depth contrast between gravity wakes (see Section 1.3) and has revealed new properties of the regions between the wakes. I also assisted Prof. Quinn with the implementation of a new symplectic method for handling Hill’s equations in a local simulation by testing his method in my local aggregates code.

6.2 Conclusions

Through the results of 100 simulations using the local cohesive aggregate code, I find that cohesion is consistent with observations ($\alpha = -2.75$ to -3 ; $R_{max} = 10$ – 20 m) in both the A ring and B rings. I have constrained the merge limit in the A ring to 0.1 – 0.5 mm/s, which brackets the Hatzes et al. (1991) result that frosty ice sticks at impact speeds below 0.3 mm/s. I have constrained the fragmentation limit in the A ring to 0.4 – 0.7 mm/s, which has no experimental reference.

The same observational criteria constrain the merge limit in the B ring to 0.25 – 0.4 mm/s, and the fragmentation limit to 0.9 – 1.1 mm/s. Alternatively, if I instead apply the cohesion parameters found for the A ring to the B ring, I find steeper size distributions ($\alpha \sim -5$ to -4) and smaller largest bodies ($R_{max} \sim 5$ – 10 m).

I discussed possible conclusions we could draw from these separate constraints in

Section 4.6. For instance, it could be that aggregates in the B ring somehow have stronger cohesive bonds through differences in the particles' surface texture there. Or perhaps the observations of the rings I am using as criteria are sensitive to different structures than is my analysis, and the size distribution of particles in the B ring truly is steeper than has been reported by observations.

I find a critical strength in both rings of 10^{-2} Pa (for $\beta = 2$), above which collisions dominate the state of the aggregate population, and stresses like tides and spin are negligible. For the sticking model presented in Hatzes et al. (1991), I estimate $\beta = 10^{-3}$, making the critical strength 10 Pa for that mode of cohesion. This value is consistent with (or lower than) the upper limits of cometary strengths presented in Section 3.1 (~ 5 Pa for SL9, ~ 100 Pa for Brooks 2, and $\sim 10^{3-4}$ for Temple 1), so it's possible that ice in the rings, were it to form comet-like frosty aggregates, could plausibly be as strong as (or stronger than) the critical strength we have found.

Applying both the discovery of a plausible critical strength and my simple scaling relationship between impact pressure and speed (see Appendix A.2), I conclude that tides and reasonable aggregate spin rates are not significant fragmentation factors for weakly cohesive aggregates in the rings, and are not the limiting factor in aggregate growth. If the strength is above the critical strength, impact fragmentation is the dominant breakage mechanism. If the strength is lower than the critical strength, then the scaling relationship (Eq. (A.8)) implies that the fragmentation limit must be so low as to make aggregate survival impossible.

The maximum size of the aggregates formed in my model is determined by the choice of fragmentation limit (so long as the strength is set above the critical value). As the plausible range of fragmentation limits is restricted by observations of α , the largest bodies this model can produce are approximately 40 m in radius. This gives evidence that “pro-pellers,” as well as Tremaine (2003)'s “shear-free” regions, are not collections of weakly

cohesive frosty ice particles.

It might be possible to provide constraints on cohesion parameters in the B ring by searching for abnormally low physical optical depths, but these variations should only occur at the largest of cohesion parameters that allow the aggregates to grow beyond 20–25 m in radius. It appears that cohesion parameters have no effect on the optical depth in the A ring.

My results are generally consistent with frosty ice cohesion occurring in the rings—though I fully acknowledge that consistency by no means implies proof. To find proof, the cohesive model could be used to explore other ring phenomena. For instance, one might determine if any relationship exists between the azimuthal brightness asymmetry and cohesion; or investigate the effects of cohesion on ring viscosity, as applied to wave damping, or overstability. Finally, true proof should be possible if the rings could be imaged with sub-meter resolution, so that we may directly resolve any aggregates of coherent particles forming in the wakes—but such a mission is not currently planned.

It is important to note that these results are all specific to the A and B rings of Saturn, and that other systems with different properties (e.g., material densities, orbital frequencies, and velocity dispersions) should respond differently to the range of cohesion parameters that match the observations of the rings.

Appendix A

Derivations

A.1 Contact Area for Intersecting Spheres

Here I 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.)

I refer to Fig. A.1 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. I 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 the right triangle, the Pythagorean theorem gives:

$$(R' + 0.5d)^2 + h^2 = R^2 \tag{A.1}$$

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

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

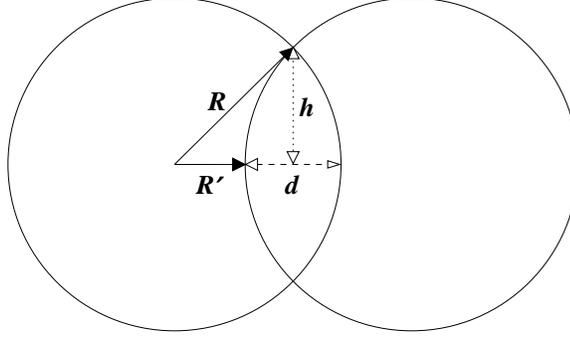


Figure A.1: Geometry and symbols used in the derivation of an approximate value for the contact area between intersecting spheres (Appendix A.1). Here, d refers to the dashed line, indicating the depth of particle overlap; h is the dotted line, indicating the distance from the center of the overlap to 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, I define a new parameter β :

$$\beta \equiv d/R \tag{A.3}$$

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:

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

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

$$\pi h^2 = \pi R^2 \beta \tag{A.5}$$

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 $\beta = 0$, there is no overlap or deformation, and the contact area is zero.

Contact area scales inversely with strength (with constant force). As such, Eq. (A.5) 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. (A.4)), 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, the prevailing bonding concept in this dissertation 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 $d \approx 25 \mu\text{m}$. With $R = 2.5 \text{ cm}$, I 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.

A.2 A Simple Model Relating Impact Pressure and Speed

Here I 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 provides understanding into 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). Thus, if desired, any source of stress that breaks an aggregate's bond can be self-consistent with any other, be it collisional or tidal.

I 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 A.1.

I 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:

$$m\Delta v = F\Delta t \quad (\text{A.6})$$

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—I am merely aiming for an order-of-magnitude estimate.

Then I 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 I 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 , I use the result of Appendix A.1, which was 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 $m = \frac{4}{3}\pi R^3\rho$, I now have:

$$\frac{4\pi R^3\rho v_{imp}}{3\pi R^2\beta(1 - \beta/4)\Delta t} \sim \frac{R\rho v_{imp}}{\beta(1 - \beta/4)\Delta t} = P_{imp} \quad (\text{A.7})$$

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

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

$$\left(\frac{1}{\beta}\right) \left(\frac{R}{1\text{ m}}\right) \left(\frac{\rho}{1\text{ g/cm}^3}\right) \left(\frac{v_{imp}}{1\text{ mm/s}}\right) \left(\frac{1\text{ s}}{\Delta t}\right) = \left(\frac{P_{imp}}{1\text{ Pa}}\right) \quad (\text{A.8})$$

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

For this example, I choose $v_{imp} = 0.5$ mm/s—approximately the escape speed of a sphere with $R = 1$ m, and $\rho = 0.5$ g/cm³. (Note that this speed is the fragmentation limit used throughout my simulations.)

One of the larger uncertainties among the quantities in Eq. (A.8) 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, I 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 A.1, I 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—the fiducial choice of strength in this project.

Bibliography

- Albers, N., Spahn, F., 2006. The influence of particle adhesion on the stability of agglomerates in Saturn's ring. *Icarus* 181, 292-301.
- Asphaug, E., Benz, W., 1996. Size, density, and structure of Comet Shoemaker-Levy 9 inferred from the physics of tidal breakup. *Icarus* 121, 225–248.
- 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.
- Barnes, J., Hut, P., 1986. A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature* 324, 446–449.
- Bridges, F.G., Hatzes, A., Lin, D.N.C., 1984. Structure, stability and evolution of Saturn's rings. *Nature* 309, 333–335..
- 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* 123, 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. and Luding, S. (Eds.), *Granular Gases*. Springer Berlin / Heidelberg, pp. 153–183.

- 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.
- 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.
- Colwell, J.E., Nicholson, P.D., Tiscareno, M.S., Murray, C.D., French, R.G., Marouf, E.A., 2009. The structure of Saturn’s rings. In: Dougherty, M.K., Esposito, L.W., Krimigis, S.M. (Eds.), Saturn from Cassini-Huygens. Dordrecht: Springer, pp. 375–412.
- 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. Dordrecht: Springer, pp. 459–510.
- Daisaka, H., Ida, S., 1999. Spatial structure and coherent motion in dense planetary rings induced by self-gravitational instability. *Earth, Planets, and Space* 51, 1195–1213.
- French, R.G., Nicholson, P.D., 2000. Saturn’s rings II: Particle sizes inferred from stellar occultation data. *Icarus* 145, 502–523.
- French, R.G., Salo, H., McGhee, C., Dones, L., 2007. HST observations of azimuthal asymmetry in Saturn’s rings. *Icarus* 189, 493–522.
- Goldreich, P., Tremaine, S., 1982. The dynamics of planetary rings. *ARA&A* 20, 249–283.
- Greenberg, R., Brahic, A. (eds.) 1984. Planetary rings. Univ. of Arizona Press.

- 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.
- Housen, K.R., Holsapple, K.A., 1999. Scale effects in strength-dominated collisions of rocky asteroids. *Icarus* 142, 21–33.
- Jacobson, R.A., Spitale, J., Porco, C.C., Beurle, K., Cooper, N.J., Evans, M.W., Murray, C.D., 2008. Revised orbits of Saturn's small inner satellites. *Astron. J.* 135, 261–263.
- 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.
- Julian, W.H., Toomre, A., 1966. Non-axisymmetric responses of differentially rotating disks of stars. *ApJ* 146, 810–830
- Karjalainen, R., Salo, H., 2004. Gravitational accretion of particles in Saturn's rings. *Icarus* 172, 328–348.
- Karjalainen, R., 2007. Aggregate impacts in Saturn's rings. *Icarus* 189, 523–537.
- 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.

- Monaghan, J.J., 1992. Smoothed particle hydrodynamics. *ARA&A* 30, 543–574.
- Murray, C.D., Dermott, S.F., 1999. *Solar System dynamics*. Cambridge University Press, Cambridge.
- Perrine, R.P., Richardson, D.C., 2007. Numerical studies of satellite-ring interactions. *AAS/Division for Planetary Sciences Meeting Abstracts*, 39, #10.01.
- Perrine, R.P., Richardson, D.C., Scheeres, D.J., 2011. A numerical model of cohesion in planetary rings. *Icarus* 212, 719–735.
- Perrine, R.P., Richardson, D.C., 2011. *N*-body simulations of cohesion in dense planetary rings: A study of cohesion parameters. *Icarus*, submitted.
- Porco, C.C., Thomas, P.C., Weiss, J.W., Richardson, D.C., 2007. Saturn’s small inner satellites: Clues to their origins. *Science* 318, 1602–1607.
- 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., 1994. Tree code simulations of planetary rings. *Mon. Not. R. Astron. Soc.* 269, 493–511.
- 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, J.E., Melosh, H.J., Lisse, C.M., Carcich, B., 2007. A ballistics analysis of the Deep Impact ejecta plume: Determining Comet Tempel 1's gravity, mass, and density. *Icarus* 190, 357–390.
- Richardson, D.C., Michel, P., Walsh, K.J., Flynn, K.W., 2009. Numerical simulations of asteroids modelled as gravitational aggregates with cohesion. *Plan. Sp. 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.
- Sekanina, Z., Yeomans, D.K., 1985. Orbital motion, nucleus precession, and splitting of periodic comet Brooks 2. *Astron. J.* 90, 2335–2352.
- 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. 126 pp.
- Supulver, K.D., Bridges, F.G., Lin, D.N.C., 1995. The coefficient of restitution of ice particles in glancing collisions: Experimental results for unfrosted surfaces. *Icarus*

113, 188–199.

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., Burns, J.A., Hedman, M.M., Porco, C.C., 2008. The population of propellers in Saturn's A ring. *Astron. J.* 135, 1083–1091.

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.

Tiscareno, M.S., Burns, J.A., Sremcevic, M., Beurle, K., Hedman, M.M., Cooper, N.J., Milano, A.J., Evans, M.W., Porco, C.C., Spitale, J.N., Weiss, J.W., 2010. Physical characteristics and non-keplerian orbital motion of "propeller" moons embedded in Saturn's rings. *ApJL* 718, 92–96.

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.