Reversing Time to Find Families:

Reviewing Backwards Integration as a New Method of Family-finding in the Kuiper Belt

Nathan S. Benfell

A senior thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Bachelor of Science

Darin Ragozzine, Advisor

Department of Physics and Astronomy

Brigham Young University

ABSTRACT

Reversing Time to Find Families:
Reviewing Backwards Integration as a New Method of Family-finding in the Kuiper Belt

Nathan S. Benfell
Department of Physics and Astronomy, BYU
Bachelor of Science

The solar system is filled with collisional families, each consisting of several objects all generated though a single historical collision. There are hundreds of known familes in the asteroid belt, but only one known family in the Kuiper Belt (an icy, rocky region beyond Neptune). The age of young asteroid collisional families is often determined by using reversed simulations (i.e. backwards integration) of the solar system. This method is not used for discovering young asteroid families and is limited by unpredictable factors unique to the Asteroid Belt (e.g. the Yarkovsky Effect). The Kuiper Belt is absent of these unpredictabilities, and thus it was theorized that backwards integrations could be an advantageous method for both Kuiper Belt Object (KBO) family finding and characterization. Such integrations are ambitious and would require high accuracy over long timescales ($\sim$ billions of years). This thesis outlines work done examining the feasibility of backwards integration as a method of family-finding, and specifically delves into the associated challenges.

Keywords: Kuiper Belt, Collisional Families, Backwards Integration

ACKNOWLEDGMENTS

First and foremost I owe a Haumea-sized thank you to my parents, Steve and Jill Benfell. Thank you for teaching me by example to do the best I can do and to be the best I can be. Thank you for loving and supporting me despite my detours, u-turns, and missteps. Thank you for the prayers before every test. Thank you for the long phone conversations about school, life, and everything in between. Thank you for teaching me who I am and what I represent.

I also owe a huge thank you to my brothers, Spencer and Thomas Benfell. Thank you for putting up with me for so long. Thank you for loving me even though you lived with me. Thank you for teaching your big brother how to be a big brother, and beyond that to be a good person. Thank you for being my best friends. #waterbuffalo.

Thank you to my roommates, friends, and classmates for getting me through tough classes and taking me on epic adventures. Thank you especially to the Astrow™ for staying up all night in pursuit of both A's and A-type stars.

And finally, a massive thank you to my advisor Dr. Darin Ragozzine. Thank you for teaching me to do astronomy and being patient with me while I found my family-finding-footing. Thank you for infusing me with a love of the outer solar system, orbital dynamics, and Haumea that will not soon fade.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Study of the football-shaped dwarf planet Haumea in the outer solar system has yielded several noteworthy findings since its discovery in 2004 (Brown et al. 2005; Santos-Sanz et al. 2005). Haumea is one of many objects that reside in the Kuiper belt, a region beyond the orbit of Neptune inhabited by several rocky, icy bodies known as Kuiper belt objects, or KBOs (the most famous of these KBOs is the beloved Pluto, itself relegated to the same dwarf-planet status as Haumea due in part to its home in the Kuiper belt). Haumea is the fastest-spinning known object in the solar system, completing one rotation on its axis in just under four hours (Lacerda et al. 2008). Perhaps as a result of this rapid spin, Haumea is highly elongated in one direction, producing a characteristic squash or pill shape (Lockwood et al. 2014; Ortiz et al. 2017). Just recently, astronomers at an observatory in Spain discovered a ring system around Haumea (Ortiz et al. 2017).

Of particular interest to dynamicists is the fact that Haumea has a collisional family, composed of itself and over two dozen other KBOs (Brown et al. 2007; Ragozzine & Brown 2007). A collisional family is formed when two large objects crash into one another, resulting in the creation of several smaller bodies with very similar orbits. The orbits of these objects (and any object on any orbit in the universe) can be defined by a set of six orbital elements, including parameters such as the size of the orbit (known as the semi-major axis), elongation of the orbit (known as eccentricity), or tilt of

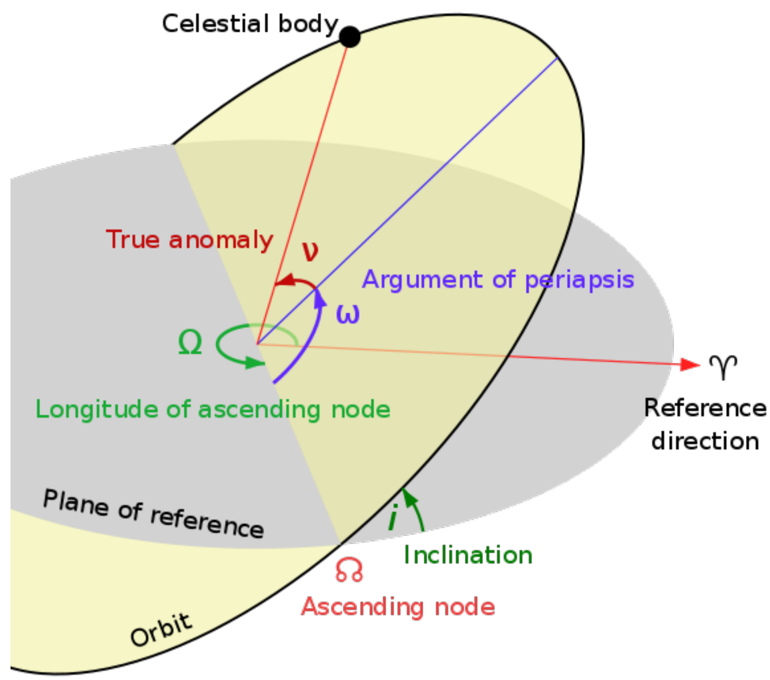

Symbol	Orbital Element	Meaning
a	Semi-major Axis	Size of the orbit
e	Eccentricity	Elliptical-ness of orbit
i	Inclination	Tilt of the orbit relative to reference plane
Ω	Longitude of the Ascending Node	Rotation of orbit about pole of reference plane
ω	Argument of Periapsis	Rotation of orbit about central axis
ν	True Anomaly	Location of body along orbital path

Figure 1.1 Description of the six orbital elements, and diagram of the elements corresponding to orientation angles (i, Ω, ω, ν). The proper elements of an orbiting body can be derived by time averaging each of its orbital elements. Image courtesy Wikipedia.

the orbit (known as inclination), along with a few orientation angles (see figure 1.1). These orbital elements can be averaged over time and thus become proper elements. The objects in families tend to cluster in one or more of these proper elements, and many families in the solar system have been discovered by searching for these clusters.

Collisional families are not uncommon, but nearly all the known families in the solar system reside in the asteroid belt (a belt of rocky, metallic objects between Mars and Jupiter). Indeed, the Haumea Family is the only currently known family in the Kuiper belt. Though there are presumed to be many more (Barr & Schwamb 2016; Levison et al. 2008), families in the Kuiper belt are impossible to find by modern asteroid family-finding methods (Marcus et al. 2011).

Yet another challenge in collisional family research is determining the age of known families. A technique commonly used to find the age of young asteroid families is that of backwards or reverse integration (Carruba et al. 2018a;b; Nesvorný & Bottke 2004; Tsirvoulis 2019). This technique involves incrementally and numerically solving the set of differential equations governing the motion of bodies in gravitational fields. Effectively, reverse integration retraces the path of celestial bodies backwards through time. Backwards integration can reveal the moment in time in which the family-forming collision occurred, manifest by the orbital elements of several objects converging sharply at the time of impact (see figure 2.2).

But the possibility of using backwards integration to *discover* new collisional families is one that has never been properly explored. As stated, backwards integration is commonly used to estimate the age of known asteroid families. However, backwards integration is not used to discover new asteroid families primarily due to dynamical effects particularly prevalent in the asteroid belt (Bottke et al. 2006), as well as the effectiveness of other modern asteroid family-finding techniques such as proper element clustering. But these dynamical effects are absent in the Kuiper belt, and proper element clustering fails to find new KBO families, suggesting the concept of KBO family-finding through backwards integrations is certainly worth exploring (Marcus et al. 2011). Though this idea is ambitious and there would certainly be challenges, it is reasonable to assume that backwards integration could open the door to discovering the many families hidden throughout the Kuiper belt. Such discoveries would have the potential to significantly advance research into the formation of the solar system and the evolution of its structure.

This notion served as the motivation for research conducted during the past two years, the results of which are outlined in this thesis. What follows is first a review of the literature on collisional family research throughout the solar system, then a detailed explanation of the experimental methods used in this project, and finally a summary and analysis of the results of the project.

Chapter 2

Background

2.1 Collisional Families in the Asteroid belt

In 1918, the Japanese astronomer Kiyotsugu Hirayama was the first to notice that several groups of asteroids seemed to share strangely similar orbits and coined the term “family" to describe these groups (Hirayama 1918). But not until the late 20th century did the famous Gerard Kuiper (after whom the Kuiper belt is named) become the first to posit that these families could have been formed through collisions (Kuiper 1974).

2.1.1 Discovery of Asteroid Families

The techniques used to discover asteroid families vary depending on the age of the family. Young asteroid families can be found through clustering in orbital element space. As stated in chapter 1 and shown in figure 1.1, any Keplerian orbit can be defined by a set of 6 orbital elements: semi-major axis (a), eccentricity (e), inclination (i), longitude of the ascending node (Ω), argument of periapse (ω), and mean anomaly (ν). Since asteroid families are formed via collision, the orbits of the remnant family members are all initially very similar. Therefore relatively recently-formed famililes

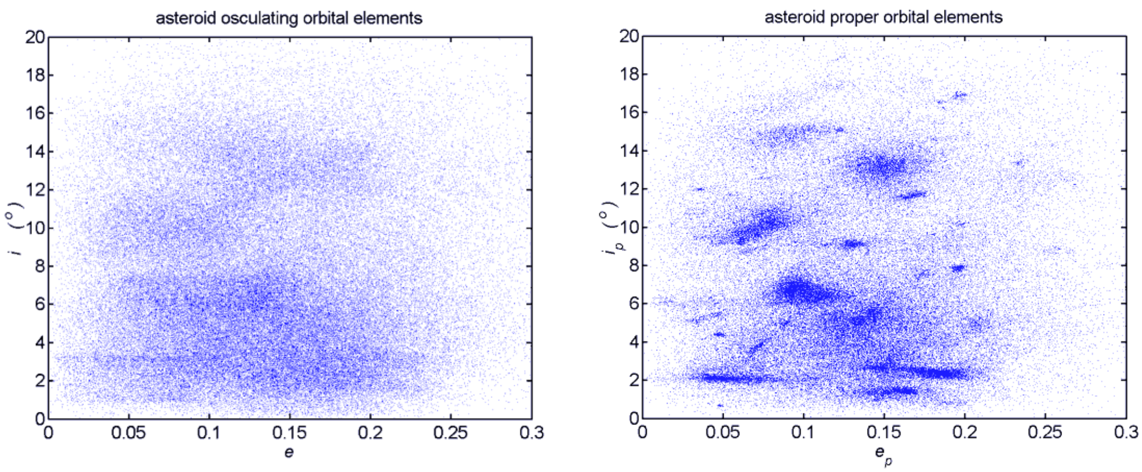

Figure 2.1 The known asteroids plotted in osculating orbital element space versus the same asteroids plotted in proper element space. In this case eccentricity (e/e_p) is plotted against inclination (i/i_p). Clusters (i.e. families) become much more frequent and apparent in proper element space, where gravitational perturbations and chaos do not diffuse orbits. Image courtesy Piotr Deuar.

(age $\leq$ 1 million years) can be found by searching for clusters in orbital element space (Bowell et al. 1994). A prime example of this methodology is the discovery of the Datura family, a family whose age is only 0.53 million years (Nesvorny et al. 2006; Vokrouhlický et al. 2009).

But families older than $\sim$ 1 million years can only be found by searching for clusters in *proper* element space (see figure 2.1). Standard (also called osculating) orbital elements are, over time, perturbed by gravitational forces from surrounding bodies (e.g nearby planets) and also subject to chaotic interactions. Thus on timescales longer than $\sim$ 1 million years, osculating orbital elements of an asteroid family are not similar enough to see clusters in element space. Proper elements are found by applying secular perturbation theory to the asteroid's orbits, resulting in essentially a time-averaged value not subject to periodic oscillations (Knezevic et al. 2002). Since proper elements remain constant over long intervals, clusters in proper element space usually indicate family members, each generated from the same collision.

Though debate exists over the most effectual means of identifying clusters, proper element clustering has led to the discovery of over 100 asteroid families to the present day. While clustering in proper element space is often visually apparent, determining appropriate cluster boundaries and identifying family members from background asteroids is more challenging. Several statistical algorithms have been developed explicitly for this purpose (Zappala et al. 1990; 1994). The increasing catalog of asteroids has bolstered both the number of families and the size of known families but has also increased the difficulty of properly identifying and differentiating clusters (Nesvorný et al. 2015). But as more and more clusters are identified, defined, and declared families, further work is being done to characterize these newfound families.

2.1.2 Age Estimation of Asteroid Families

The chief component of asteroid family characterization is achieving an estimate for the age of the family (t_{age}). Family ages in the asteroid belt can vary from hundreds of thousands of years to multiple billions of years (Nesvorný et al. 2015; Nesvorny et al. 2006; Vokrouhlický et al. 2009). This wide range of family ages necessitates a varied approach to age determination based on the family in question.

Relatively young families ($t_{age} \leq 10$ million years) can be aged through the use of backwards integrations. By numerically and iteratively solving gravitational differential equations for the bodies in the family, the precession of the orbits through time can be reversed until the point of collision (i.e. t_{age}). In the time immediately following the collision, the newly created family members are tightly confined to a specific region in proper element space, specifically in Ω (the longitude of the ascending node) and ω (the argument of periapse), the elements that define the rotational orientation of the orbital plane. Thus as the backwards integrations approach t_{age}, the values of Ω and ω for the asteroid family members sharply converge, pinpointing the age of the family (see figure 2.2). This technique was pioneered on the Karin asteroid family by researchers David Nesvorny and William Bottke of the Southwest Research Institute and has since been used to

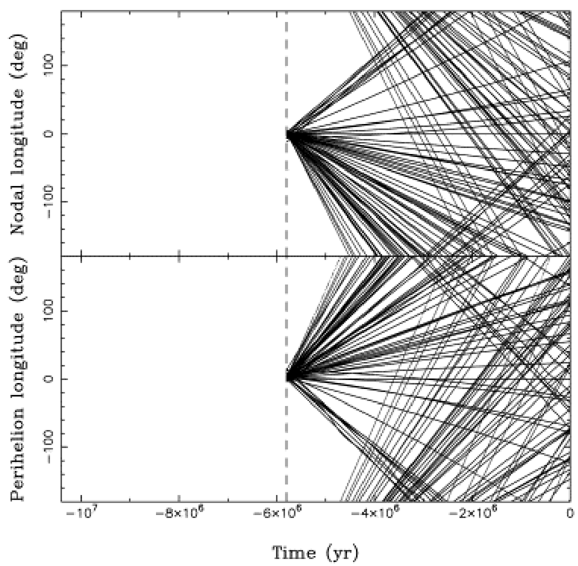

Figure 2.2 The results of a backwards integration of the Karin family. The proper elements Ω (longitude of the ascending node) and ω (argument of periapse) converge tightly at $t_{age} = 5.8 \pm 0.05$ million years. Image courtesy Nesvorný & Bottke (2004).

characterize the age of several other young asteroid families (Carruba et al. 2018a;b; Nesvorný et al. 2002; Nesvorný & Bottke 2004; Tsirvoulis 2019).

However, the use of backwards integration as a method of family age-estimation in the asteroid belt is limited to use on young families by dynamical influences unique to the asteroid belt. This limitation is imposed primarily by the Yarkovsky effect. This effect refers to the impact of radiative heating on the motion of bodies in the solar system. Specifically, the delay between when photons emitted from the sun hit the surface of a body and that body increasing in temperature causes a differential force in the direction of motion of the object. Over time this delay increases the orbital velocity, which in turn increases the semi-major axis (a) of the body. This effect is minimal for large bodies but can be prominent for bodies with diameters ≤ 20 kilometers (Bottke et al. 2001). The asteroid belt — and therefore any asteroid family — is filled with these small bodies, and thus on timescales longer than 10 million years, the dynamical spreading of family members induced by the Yarkovsky effect inhibits the effectiveness of backwards integration for age estimation. To date the oldest family successfully backwards integrated to t_{age} is the Veritas family, of age 8.3 ± 0.5 million years (Nesvorný et al. 2003; Tsiganis et al. 2007). Thus, for older families an alternate age estimation mechanism is required.

Indeed, for older families the gradual dynamical spreading of the Yarkovsky effect is advantageous. Typical timescales for the Yarkovsky drift of asteroids are known and understood (Bottke et al. 2006; 2001; Nesvorný et al. 2015). As stated, the Yarkovsky effect more prominently influences asteroids of smaller mass. But the Yarkovsky effect also influences objects of lower albedo (i.e. brightness) more heavily. Bright objects reflect photonic energy and therefore resist changes in temperature, reducing the Yarkovsky effect. Conversely, dark objects absorb photons and experience greater temperature changes, increasing the Yarkovsky effect. By combining knowledge of the Yarkovsky drift timescales with data on the distance (a), density (ρ), and albedo (p_V) of an asteroid family, a rough age estimate may be attained from the equation:

$$t_{age} \simeq 1 \text{ g.y.} \times \left(\frac{C_0}{10^{-4}AU}\right) \times \left(\frac{a}{2.5AU}\right)^2 \times \left(\frac{\rho}{2.5}\right) \times \left(\frac{0.2}{p_V}\right)^{1/2}$$

where g.y. identifies the units (billions of years) and C_0 is a constant associated with the drift timescale of asteroids due to the Yarkovsky effect (Nesvorný et al. 2015). Using this method, researchers were able to estimate ages of families as old as 1.3 ± 0.5 billion years, with the caveat that uncertainties grow as age increases (Carruba & Morbidelli 2011; Vokrouhlický et al. 2006).

All told there is a wealth of current knowledge on collisional asteroid families and multiple means of discovery and age estimation. Additionally, the relatively close proximity of the asteroid belt to earth has resulted in a large number of known asteroids and small uncertainties in their characteristics and orbital parameters, contributing to both a high quantity of families and of family members. The same cannot be said, however, for collisional families in the Kuiper belt.

2.2 Haumea and Collisional Families in the Kuiper belt

In 2003, the observational astronomer Mike Brown of the California Institute of Technology found what was the latest in a series of new dwarf planets discovered in the Kuiper belt in the early 2000s: Haumea (Brown et al. 2005; Santos-Sanz et al. 2005). Unlike other bodies in the solar system of comparable or larger mass (which have gravitationally pulled themselves into a spherical shape), Haumea's shape is that of a triaxial ellipsoid (i.e. a football) (Rabinowitz et al. 2006). This is presumed to be a result of its extremely fast rotation period, as Haumea spins once on its axis in just four hours (Lacerda et al. 2008; Rabinowitz et al. 2006). Haumea is known to have two satellites: the outer, brighter moon Hi'iaka and the inner, dimmer moon Namaka (Ragozzine & Brown 2009). Surrounding Haumea is a faint ring, discovered recently as the body passed in front of a distant star (Ortiz et al. 2017; Winter et al. 2019).

Of principal importance to this review, Haumea is also unique among KBOs for its collisional family, the only currently known family in the Kuiper belt (Brown et al. 2007; Ragozzine & Brown 2007). While the discovery of the Haumea family confirmed the existence of collisional families in the outer solar system, it also served to highlight the difficulties of family research in the Kuiper belt.

2.2.1 Discovery and Characterization of the Haumea Family

The Haumea family was discovered in a manner entirely different from the typical discovery of asteroid families discussed in section 2.1, namely through spectroscopic analysis. An overview of the surface composition of any extraterrestrial object can be attained through a process known as spectroscopy. This technique takes advantage of the differences in which light interacts with various molecules. Any object will emit a unique spectra, in essence a light-generated fingerprint that details exactly what compounds cover its surface. Upon examining the infra-red spectra of several KBOs, a team of astronomers led by the aforementioned Mike Brown discovered several, including Haumea, with consistently anomalous compositions. (Brown et al. 2007). Specifically, while most KBOs are covered by layers of methane ice, these KBOs were completely absent of methane, instead covered with water ice. Further analysis of the orbits of these objects confirmed that they also all inhabited a tight region in proper element space, confirming their identity as a collisional family (Brown et al. 2007; Ragozzine & Brown 2007).

Since the discovery of the Haumea family, much work has been devoted to its characterization, particularly the family's age and means of formation. The current age estimate of 3.5 ± 2 billion years was made by Ragozzine & Brown (2007). This estimate was made by forwards-integrating several different simulated KBO objects with assumed initial eccentricities and finding the integration time necessary to create the eccentricity diffusion present in the modern Haumea family. But the large error bar of ± 2 billion years combined with the young age of the solar system of 4.5 billion years effectively limits the origination of the Haumea family as merely being primordial (i.e. between the formation of the solar system and 1 billion years ago). While subsequent studies have confirmed this estimate, they have failed to refine the uncertainty (Volk & Malhotra 2012).

Precise details also remain elusive regarding the formation of the Haumea family. Several theories have been put forth, with formation mechanisms ranging from single catastrophic collisions to long-term rotational fission causing small fragments of Haumea to fly off of its fast-spinning ends (Kondratyev 2016; Leinhardt et al. 2010; Levison et al. 2008; Ortiz et al. 2012; Ragozzine & Brown 2007). Increasingly complex theories have come to the fore in recent years, hypothesizing family formation through a sequence of collisions (such as the model shown in figure 2.3) or mergers (Leinhardt et al. 2010; Schlichting & Sari 2009). But these theories are not without issues, and none is able to sufficiently account for all of the current parameters of the family (Campo Bagatin et al. 2016; Proudfoot & Ragozzine 2019). As a result, the formation of the Haumea remains a mystery.

In short, despite a large body of research on the Haumea family since its discovery in 2007,

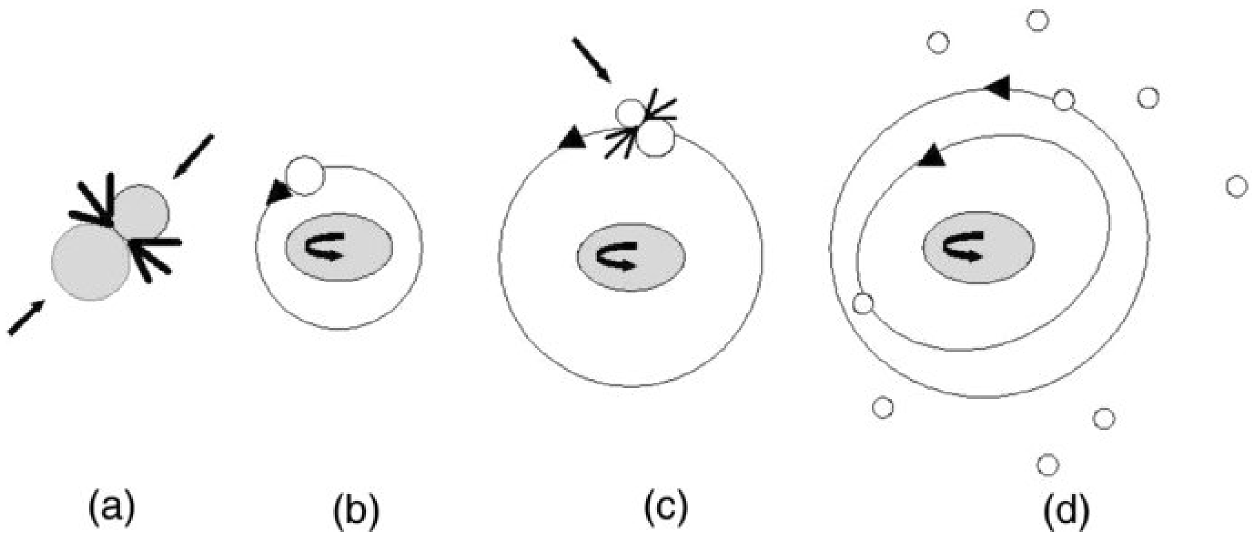

Figure 2.3 Diagram of a Haumea family formation model proposed by Schlichting & Sari (2009). In part (a) a destructive collision occurs between Haumea and another large object. Following the collision the debris accretes into a single satellite. In part (b) the satellite experiences tidal forces that gradually increase its semi-major axis. Next, in part (c) the satellite undergoes a destructive collision. Finally, in part (d) the resulting debris scatters to form the Haumea family and the moons Hi'iaka and Namaka. Image courtesy Schlichting & Sari (2009).

much remains unknown about both its age and formation. Notwithstanding the many uncertainties around the family, the mere fact of its existence suggests that there are more KBO families awaiting discovery. However, family-finding efforts in the Kuiper belt take a much different form than the previously discussed family-finding efforts in the asteroid belt.

2.2.2 Kuiper belt vs. Asteroid belt Family Research

Though the Haumea family remains the only known Kuiper belt family, many more are likely present (Barr & Schwamb 2016; Chiang 2002; Levison et al. 2008; Rabinowitz et al. 2011). But a number of differences exist between the dynamics of asteroid families and those of KBO families. First, the great distance and dim nature of the Kuiper belt means that many observation-derived values for KBO orbital elements are highly uncertain. This is manifest in the uncertainty of the current Haumea family age estimation mentioned in the previous section. However, new sky surveys and telescope missions incrementally increase the accuracy of these measurements, so orbital uncertainties tend to consistently decrease over time (Bannister et al. 2016).

Second, known bodies in the Kuiper belt tend to be much larger than known asteroids. Therefore collisions between these objects are much larger and result in much higher velocity dispersions than their asteroid counterparts (Ragozzine & Brown 2007). This means that the spread in orbital and proper elements among KBO family members will be much higher than the spread among asteroid families. As a result, the typical methods of orbital/proper element clustering used to identify asteroid families are ineffective in the Kuiper belt (Chiang 2002; Ragozzine & Brown 2007). In fact, researchers at the Harvard-Smithsonian Center for Astrophysics quantitatively showed that KBO families are undifferentiable from the background KBO population in proper element space (Marcus et al. 2011).

Additionally, for the same reason that proper element clustering is an invalid family-finding technique in the Kuiper belt (large size of KBOs), the Yarkovsky effect on known KBOs is largely negligible (Marcus et al. 2011; Ragozzine & Brown 2007). While this increases the long-term stability of the Kuiper belt, it negates the use of Yarkovsky dynamics for age determination described in section 2.2.

However, while the Kuiper belt is in many ways a more challenging location to find collisional families, it also has an advantage over the asteroid belt. As mentioned, the Kuiper belt is more dynamically stable in the long-term than the asteroid belt (Celletti et al. 2007; Kuchner et al. 2002; Lykawka 2012; Ragozzine & Brown 2007). Principally, this increases the timescales over which backwards integrations of KBOs can remain viable and accurate. So while asteroid families and KBO families share many similarities, enough differences persist to make family-finding and family-characterization attempts completely different between the two regions.

2.3 Gaps & Areas of Further Research

It is clear that many areas exist in which further research could lead to new and exciting results. New asteroid families continue to be discovered each year, and hundreds of newly discovered asteroids are added to existing families as well. Further refinement of asteroid family clustering algorithms could help filter through the ever increasing catalog of asteroids and resolve the issue of increasing numbers of overlapping families (Nesvorný et al. 2015; Zappala et al. 1990; 1994).

New and updated observation missions constantly return orbital information on new objects and refined information on existing objects. Re-assessment of old research with new data could lead to several discoveries and modifications to current knowledge. This could be particularly useful in the study of the Haumea family. Multiple studies suggest that more family members could lead to a more precise age estimation (Ragozzine & Brown 2007; Volk & Malhotra 2012). It is also likely that new data (and a refined age estimate) could shed more light on the formation of the family and either validate a current hypothesis or give rise to a completely new theory (Campo Bagatin et al. 2016; Proudfoot & Ragozzine 2019).

But perhaps the area with the greatest promise for discovery is in finding new collisional families in the Kuiper belt. Importantly, the potential benefit of such an endeavor is huge. As outlined by Marcus et al. (2011): "[KBO Families] provide a unique testbed for theories of the dynamical, collisional, interior, and surface properties of KBOs. If found, further KBO families will provide valuable insights into properties of individual objects and the outer solar system as a whole that are otherwise difficult to obtain."

However, as outlined in section 2.2.2, current family finding methods for asteroid families fall short in the Kuiper belt for a variety of reasons. But due to the long-term dynamic stability of most KBOs, a possible new opportunity exists for KBO family-finding. While backwards integrations are only tenable to lengths $\leq$ 10 million years in the asteroid belt, the same is not true in the Kuiper belt. In fact, due to the high level of KBO stability it is reasonable to assume that backwards integration could serve as a viable family-finding method in the Kuiper belt. While these integrations would need to maintain accuracy over long timescales, increased orbital data accuracy from new missions, coupled with new developments in bit-by-bit reproducible integrators add to the efficacy of this proposal (Karlsson 2011; Rein & Liu 2012). Such an endeavor could lead to identifying families as visually striking as the age characterization of the Karin family shown in Figure 2.2, but more importantly could enable the discovery of many as-yet-unknown KBO families.

Knowledge of more KBO families would greatly advance current understanding of the evolution of the entire outer solar system. Study of collisional families naturally increases understanding of collisions. Greater knowledge of the collisional history of the Trans-Neptunian region would provide much insight into the formation of the outer ice giants Uranus and Neptune (suspected to have accreted mass through collisions). Additionally, knowledge that KBOs in a family share the same progenitor would allow spectral analysis of surface compositions of the family members to reveal the varying effects of weathering and orbit on the surfaces of different KBOs. Such advances would drastically impact current planetary science research, which is in large part concerned with extraterristrial geological evolution and weathering activity (Marcus et al. 2011). In sum, the discovery of new KBO families would cause ripples throughout the field of solar system astronomy. Backwards integration is an as-yet untested method of KBO family finding, and therefore deserves to be further explored.

Chapter 3

Methods and Results

The main goal of this research was to determine the feasibility of backwards integration as a family-finding method in the Kuiper belt. This chapter begins with a section describing the software used to carry out these integrations. The following three sections then describe the specific challenges — observational uncertainty, systematic uncertainty, and numerical chaos — associated with backwards integrations, the research conducted into the relative impact of these challenges, and the results of this research and their repercussions on the feasibility of backwards integration as a family-finding method

3.1 Symplectic Integrators and the REBOUND Software Package

3.1.1 Overview of Symplectic Integrators

Tools commonly employed by orbital dynamicists are symplectic integrators. At a basic level, symplectic integration consists of iteratively (through time) solving the Hamiltonian equations of motion for a given system:

$$\dot{p} = \frac{\partial H}{\partial q} \qquad \dot{q} = \frac{\partial H}{\partial p}$$

where q and p are the canonical coordinates of position and momentum respectively for a certain degree of freedom and H is the Hamiltonian. When applied to the case of celestial mechanics, such integrations can be used to evolve orbits over both short and long timescales.

Of particular utility to celestial dynamicists are symplectic integrators' ability to handle n-body systems. In the case of this project, the Kuiper belt objects in question experience gravitational influence from not only the sun, but also large nearby planets (e.g. Jupiter, Neptune, etc.), passing KBOs, and even faraway bodies such as the predicted Planet X. Taking even these small perturbations into account is important for meaningful results, as the viability of backwards integration as a family-finding method is contingent upon the accuracy of the integrations over large timescales. As a result, symplectic integrators' n-body capability mark them as the perfect tool for this project.

3.1.2 Overview of REBOUND Software Package

REBOUND is a software package developed by Hanno Rein of the Institute for Advanced Study that is designed to carry out integrations of celestial systems (Rein & Liu 2012). REBOUND has the capability to use several different symplectic integrators, including the Wisdom-Holman method tailored specifically for Keplerian systems such as those of Kuiper belt objects (Wisdom & Holman 1991). The package also contains an improved and more efficient version of the Wisdom-Holman algorithm known as WHFast, developed by Rein & Tamayo (2015).

REBOUND allows for the easy addition and removal of both user-created and real-life particles from integrations. Users can specify integration length, time-step size and save frequency. As integrations progress, orbital data is saved to proprietary files known as SimulationArchives (Rein & Tamayo 2017), which contain orbital elements and orbital velocities for every particle in the integration at each save interval. These orbital values can also be interpolated for times in between the save intervals. SimulationArchives can even be accessed and analyzed before the conclusion of the integration.

All of the integrations conducted for this research were run using REBOUND. Specific parameters were varied (as described in section 3.3.1), but most of the integrations were run using WHFast. All of the integrations are stored in SimulationArchives on the Haumea computer system under the directory /home/benfelln/research/BIFFF, and a description of the integrations is contained partially in table. 3.2, and in full in the readme file found at /home/benfelln/research/BIFFF/README.

3.2 Exploring Observational Uncertainty

Kuiper belt objects are both very far away and very dim, and thus observing them is highly difficult. In fact, apart from Pluto, every known KBO has only been discovered within the last 30 years as observational technology has progressed and higher fidelity telescopes have seen first light. In addition to limiting the discovery of new objects, the difficulty of KBO observations has led to high levels of uncertainty in orbital measurements of known KBOs. In the case of this research, these large uncertainties become a particular problem upon conducting long-term integrations. As objects are integrated through time, the orbital element uncertainties of these objects grow, and at some point become large enough that the object's orbit cannot be known with any degree of confidence.

3.2.1 Propagation of Orbital Uncertainties Over Time

Therefore, the first question requiring answering is: how long does it take for KBO orbital elements to become prohibitively uncertain? To address this question, individual integrations were run on over 1500 known KBOs. These integrations ran through the age of the solar system (4.5 billion years). The uncertainty in the longitude of the ascending node (Ω), was then plotted over the length of the integration, and the point in time marked at which the uncertainty in Ω crossed a >1 radian threshold. These times for every KBO were then plotted in a logarithmic histogram that can be seen

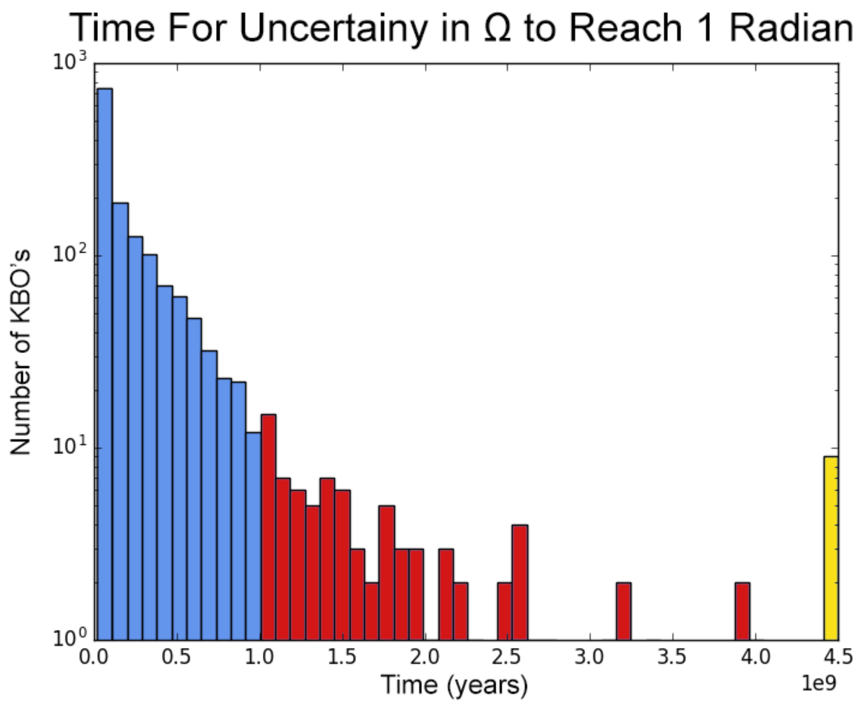

Figure 3.1 The lengths of time for each of 1511 known KBOs integrated over 4.5 Gyr to reach an uncertainty threshold of 1 radian in Ω. The red section highlights the 94 KBOs that lasted for longer than 1 Gyr, while the yellow section highlights the 9 KBOs that lasted longer than 4.5 Gyr before becoming prohibitively uncertain.

in figure 3.1.

Of the 1511 integrated KBOs, only 94 (6.2%) lasted for at least 1 billion years before reaching the 1 radian threshold and becoming prohibitively uncertain. Of those 94, only 9 (0.6%) lasted beyond the age of the solar system before reaching the uncertainty threshold. At first glance this value seems very low, and at one level it certainly is. The Haumea family is estimated to be roughly 3.5 ± 2 Gyr old, and so the prospect of using backwards integration to find families of an age on the order of the Haumea family seems implausible.

However, the numerical process used to calculate these certainty timescales is somewhat biased towards stable KBOs, disproportionately populating the leftward portion of the histogram. Additionally, even with this caveat there are a substantial number of KBOs occupying shorter certainty regimes. There are currently dozens of known asteroid families with ages of merely millions of years. Should there be families of similar age in the Kuiper belt, there are certainly enough KBOs that maintain sufficiently certain orbits to discover them through backwards integration.

Further, forthcoming data from new telescope missions will reduce the severity of this uncertainty problem. The Gaia space telescope (launched in 2013) aims to drastically refine the precision of the current catalog of celestial objects, including KBOs. Increased measurement precision will also increase the certainty timescales for known KBOs, shifting the entirety of the histogram to the right. Meanwhile, the ground-based Large Synoptic Survey Telescope (LSST, set to see first light in 2021) in Chile is poised to discover and precisely characterize about 40,000 KBOs, a factor of $\sim$20 greater than the current number of known KBOs. An increased catalog of KBOs will then serve to dramatically increase the number of KBOs with long-term certainty timescales.

Thus, while current KBO data dictates that only relatively short-term (< 1 Gyr) integrations are viable for family-finding, this is not a major stumbling block for backwards integration as a family-finding method generally. Finding a new young family would still constitute a significant scientific discovery, and future data from modern telescopes will continuously expand the limits of backwards integration.

3.2.2 Probabilistic Approach to Orbital Analysis

A more nuanced approach to analyzing integrated orbital data may also serve to increase the outer time limits of backwards integrations. This nuance consists of using probabilistic orbital analysis rather than deterministic orbital analysis. A test of this method was conducted using the Haumea family.

Specifically, integrations were constructed and run containing Haumea and every known member of the family, along with 5 clones of each body with orbital elements selected randomly from within the observational uncertainty ranges of the object. At each time step, a calculation was made of

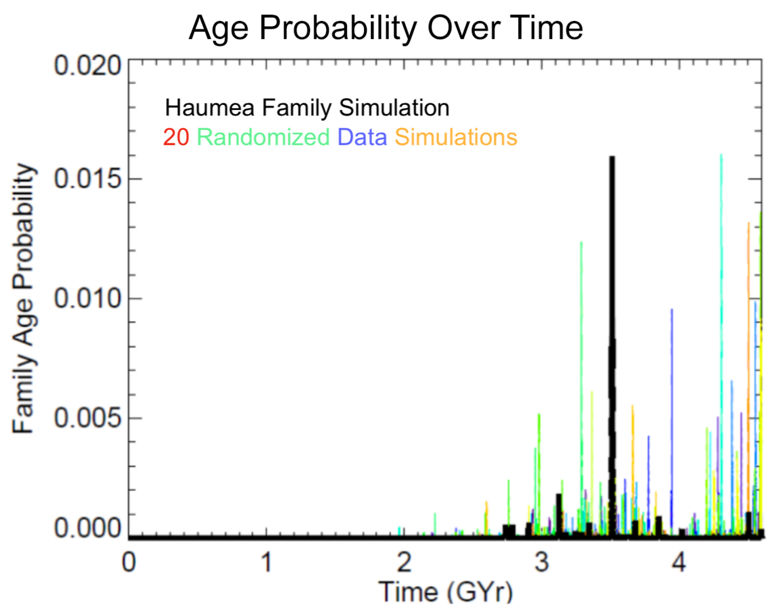

Figure 3.2 The probability estimations of the age of the Haumea family plotted over time. The black line represents a data set based on the Haumea family itself, containing 5 clones for every known family member. The colored lines represent the same probabilistic analysis carried out on randomized datasets. The spike in the original dataset is present at 3.512 ± 0.003 Gyr, aligning very well with the current age estimation of 3.5 ± 2 Gyr (Ragozzine & Brown 2007). The comparable probability spike in a single randomized dataset indicates a false alarm probability of roughly 5%.

the statistical probability (p-value) that the orbital elements of Ω and ω for a set of KBOs were identical (based on propagated uncertainty values). Plotting these probabilities over time yielded the graph shown in figure 3.2.

The distinct spike at around 3.5 Gyr indicates the most likely age of the Haumea family, and is consistent with the current age estimation by Ragozzine & Brown (2007) of 3.5 ± 2 Gyr. The outlying nature of this spike increases the confidence of this result. This same probability analysis was then carried out for 20 randomized integrations, one of which yielded a probability spike of comparable size to the original data set. This suggests a false alarm probability – the chance that a spike of this height is due purely to chance – of roughly 5%, which, while higher than desired, is low enough to suggest that the spike in the original set is likely trustworthy.

This testing yielded a specific age estimation for the Haumea family of 3.512 ± 0.003 Gyr. While much further testing would need to be carried out to precipitate the publishing of this result, it is certainly an indication of the promise of probabilistic analysis for increasing orbital certainty over long timescales. Applying this to backwards integrations designed to find families would certainly extend the uncertainty threshold well beyond the 1 Gyr line discussed in the previous section.

3.3 Exploring Systematic Uncertainty

Perhaps the most immediate concern for the viability of backwards integration as a family-finding method is that of systematic uncertainty. Specifically, this issue consists of whether integrations can be constructed that represent the solar system accurately enough to produce trustworthy results. One could theoretically run a perfect integration that accounts for every possible parameter and perturbation. But such an integration would take much too long to complete (with a limit approaching the age of the solar system). On the other hand, extremely efficient integrations (neglecting a number of parameters that may be important) likely would not be sufficiently accurate, defeating the purpose of running the integrations in the first place. Thus, the challenge remains to find optimal integration parameters that balance efficient run times with accurate results.

3.3.1 Relative Impact of Various Integration Parameters

As stated in section 3.1, the REBOUND integration package allows for several integration parameters to be specified by the user. Such options include the number and type of integrated particles, integrator method, and time step. To compare the relative importance of these various parameters, multiple integrations of a simulated KBO family based around the KBO 2002 TX300 (a large member of the Haumea family) were run, each with a single varied parameter. A single control integration was run with pre-determined "standard" parameters. A summary of the integrations can

Control Integration	
Parameter	**Value**
Integrator	WHFast
Timestep	0.25
Bodies	8 Planets Massive 2002 TX300 20 non-massive family members Planet X
Integration Settings	Yes safe_mode No corrector

Table 3.1 The parameters of the control integration. The time step is in proprietary "REBOUND-units." Safe_mode refers to a REBOUND function that deals with the propogation of Jacobi coordinates, turning it off speeds up integration time but could lead to pitfalls. Correcter refers to using an extra symplectic corrector to synchronize integration results to a higher order than the default second order.

be seen in tables 3.1 and 3.2.

Subsequent comparison of the results of these integrations allowed for analysis of the individual impact of these parameters. Example comparison plots can be seen in figures 3.3 and 3.4. Certain parameters seem to be somewhat negligible in impact. Specifically, Planet X, changing the safe_mode and corrector integration settings, integrator type, and folding the mass of the inner terrestrial planets into the Sun had little to no impact on the integrations. Even changing the time step (up to a certain point) had little impact. On the other hand, certain parameters drastically altered the results, most importantly those relating to the mass found in and around the Kuiper belt. As a result, the conclusion can be drawn that several parameters (such as Planet X) can be left out of integrations in the name of efficiency, while others (such as KBO masses) are vital to account for.

Another valuable finding from these comparisons was the discovery that all particles in the integrations seem to have inherent stability timescales. As shown in figure 3.3, some particles

Integrations Run	
Integration	**Parameter(s) Altered**
01	None
02	Timestep of 0.5
03	Timestep of 1.0
04	No safe_mode
05	Yes corrector
06	No safe_mode & yes corrector
07	4 massive family members, 16 non-massive family members
08	Mercurius integrator
09	Shift initial position of Uranus by 1 meter
10	Mercury folded into sun
11	Inner planets folded into sun
12	Mercury folded into sun & time step of 0.1
13	Inner planets folded into sun & timestep of 2.0
14	Non-massive 2002 TX300
15	time step of 0.1
16	IAS15 integrator, inner planets folded into sun, time step of 1.0
17	IAS15 integrator, inner planets folded into sun, time step of 4.0
18	IAS15 integrator, inner planets folded into sun, 4 massive family members
19	IAS15 integrator, inner planets folded into sun, 2 doubly massive family members
20	WHFast integrator, inner planets folded into sun, time step of 1.0
21	WHFast integrator, inner planets folded into sun, time step of 4.0
22	WHFast integrator, inner planets folded into sun, 4 massive family members
23	WHFast integrator, inner planets folded into sun, 2 doubly massive family members

Table 3.2 Every integration run. The first several involved changes of individual parameters. Later, several were changed at once to observe the interplay of the various parameters.

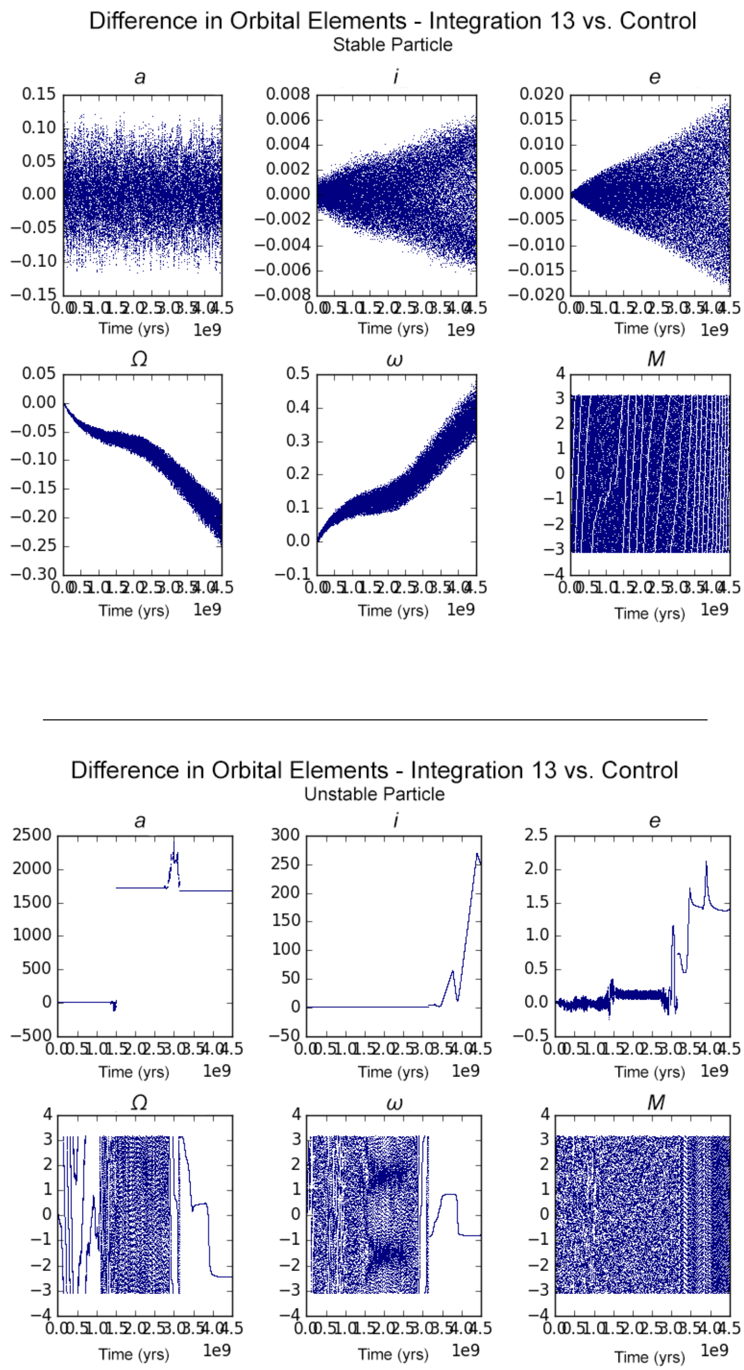

Figure 3.3 Two examples of comparisons between every orbital element for a specific particle in a given integration (in this case integration 13, in which the mass of the inner planets was folded into the sun) versus the control integration. The top set of plots is for a particle that maintained relative stability throughout the integrations, while the lower set are for a particle that was highly unstable. Comparisons such as this highlighted the relative stability timescales of various particles in the integrations.

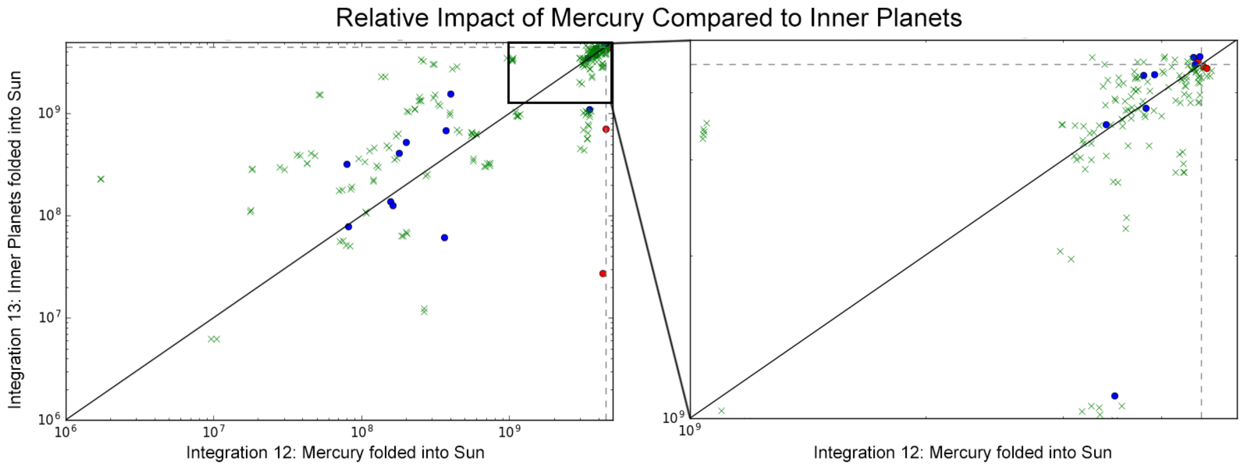

Figure 3.4 An example of a comparison between every particle in two integrations, in this case between integration 12 (in which the mass of Mercury was folded into the Sun) and integration 13 (in which the mass of the inner planets was folded into the Sun). Particles are plotted at the times at which their value of Ω became more than 1 radian different from their value of Ω in the control integration. Planets are plotted as red dots, KBOs as blue dots. The green crosses represent comparisons between pairs of KBOs. Of significance is the high concentration of objects in the upper right corner of the graph, signifying that these two chosen parameters do not make a significant difference in the orbital parameters of KBOs.

very rapidly become unstable, deviating from the control by large amounts almost immediately. By comparison, other particles remained fairly similar to their control integration counterparts across billions of years. Timescales between these two extremes were also present throughout the integrations. These timescales remained consistent for each particle for every comparison, suggesting that the particles' stability regimes were inherent, and not related to the changed parameters.

3.3.2 Outer Limit of Integrations

With these findings in mind, the next step was to test the outer limits of REBOUND integrations run with optimum parameters. A set of standard parameters was chosen based on the tests described in the previous section, and are outlined in Table 3.3. Forward integrations of a simulated family

Outer Limit Testing Integration Parameters	
Parameter	**Value**
Integrator	WHFast
Timestep	2.0
Bodies	4 Outer planets (Inner planets folded into sun) Massive 2002 TX 300 20 non-massive family members
Integration Settings	No safe_mode Yes corrector

Table 3.3 The parameters of the integrations run to test the outer limits of REBOUND integrations of a KBO family.

around 2002 TX300 were carried out over specified time intervals ranging from 5 million years to 4.5 billion years. These integrations were then run in reverse through the same amounts of time.

Should the systematic error of the integrations be low, the orbital elements of the objects in the backwards integrations should converge at times equal to the length of the corresponding forward integrations. By applying this analysis to integrations of increasing length and comparing the relative convergence in orbital element space, a reasonable estimate could be made of the maximum integration times that still preserved passable accuracy. A preliminary clustering statistic was also created to estimate more quantitatively the relative clustering of the family members following the backwards integrations. This statistic was calculated using the median average distribution of Ω across the family members, with a lower value indicating closer clustering. The results from selected integrations can be seen in figures 3.5 and 3.6.

Short-term integrations, such as that of the 10 million year old family, experienced nearly perfect convergence of family members in Ω-space, and likewise display tight clustering. Increasing the age

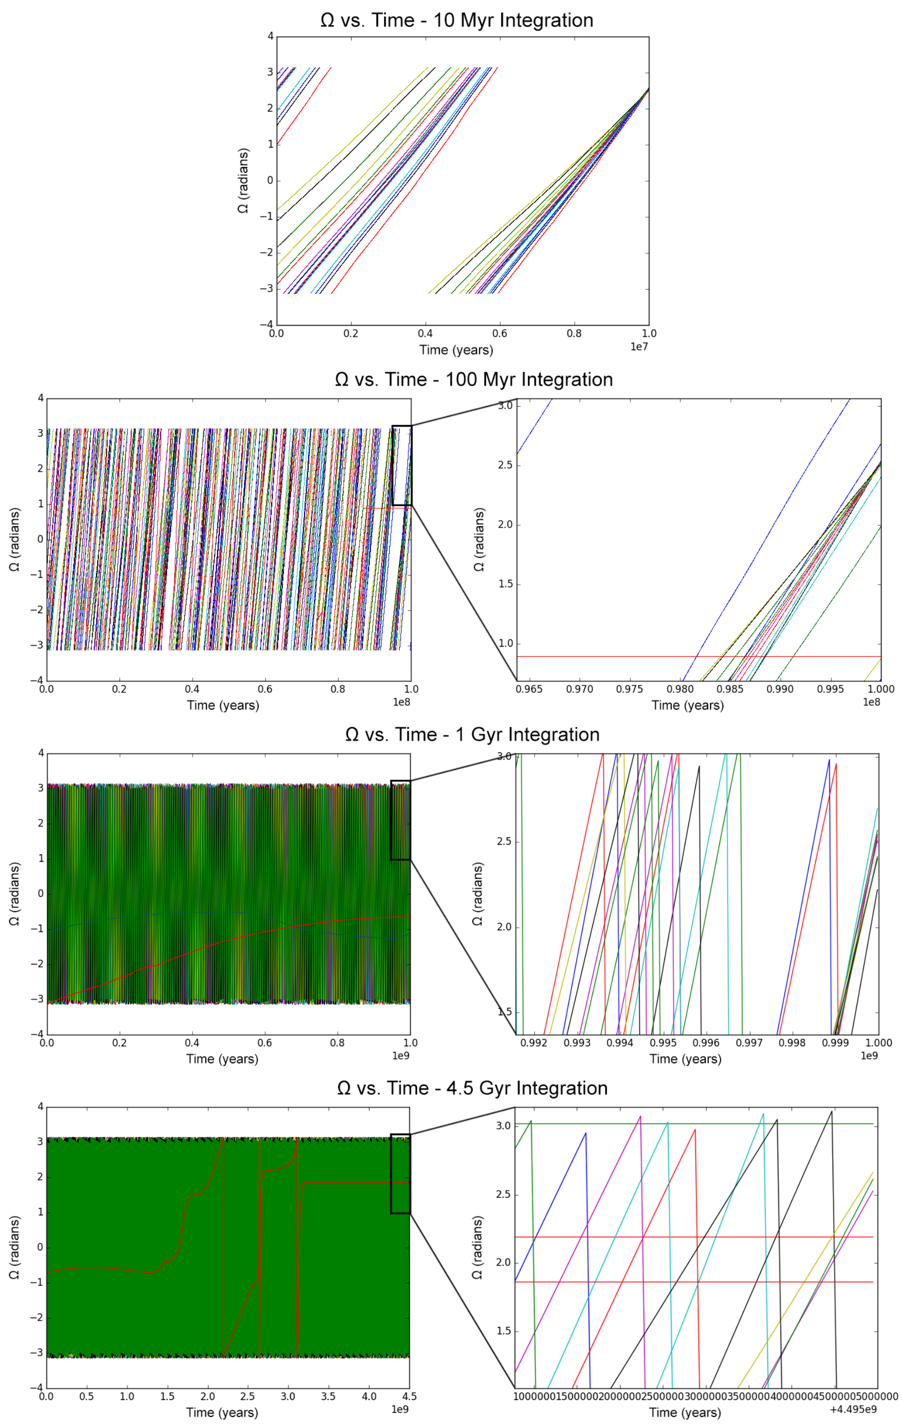

Figure 3.5 Selected examples of backwards integrations of a simulated family. In each case the age of the family is equal to the length of the integration. The shorter integrations experience almost complete convergence of the family members in Ω-space. As the family age increases, the number of converging KBOs decreases.

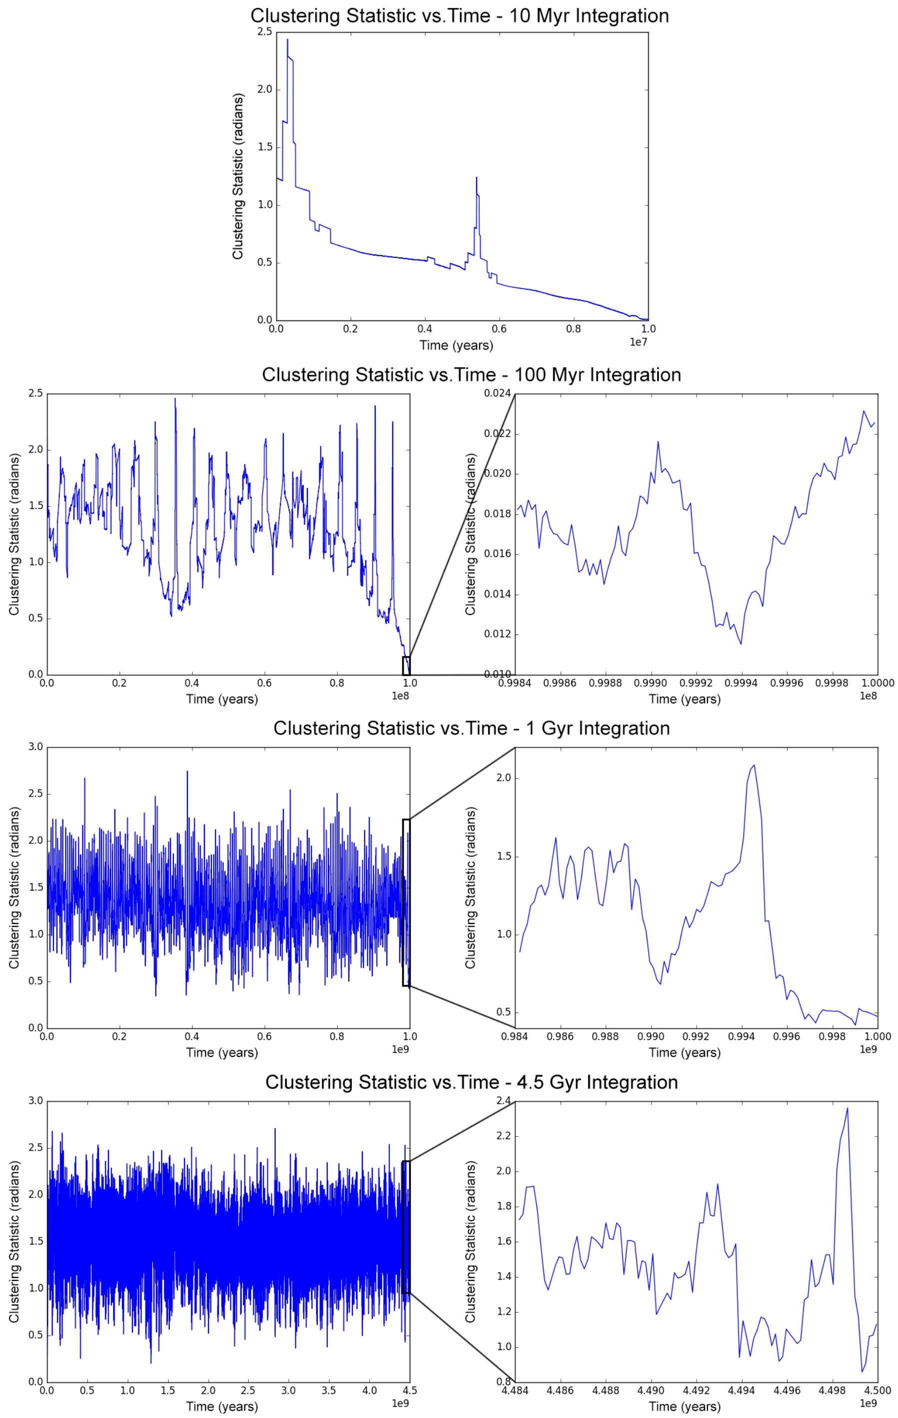

Figure 3.6 Selected examples of the clustering statistic for backwards integrations of a simulated family. The lower the statistic the tighter the cluster. In each case the age of the family is equal to the length of the integration. These results mirror the results displayed in figure 3.5. The shorter integrations show tight clustering at the family age, with clustering become less constrained as family age increases.

of the family by a single order of magnitude to 100 million years caused a slightly larger distribution of Ω among family members, but still an encouraging level of convergence that is clearly identifiable in the clustering value. This suggests that young KBO families, should they exist, would almost certainly be within the realms of detection by backwards integration. Though future work would need to be devoted to the exact methodology of family identification, the relatively high efficacy of <100 Myr backwards integrations indicates a high likelihood of success.

Even some longer term integrations such as that of the 1 billion year-old family displayed encouraging results. While more KBOs are seen to scatter from the parent family, a surprisingly large number of KBOs still converge to the expected value in Ω-space, with an ending clustering value of less a half-radian. Thus, though not as visually striking as the clustering of younger families, the convergence in element space of families as old as 1 Gyr is likely enough to yield potential family-finding results.

However, very long-term integrations display somewhat more problematic results. The integration of a 4.5 billion year old family revealed only three of the original 20 family members converging in Ω. The large number of scattered objects suggests that, at least with the chosen parameters, backwards integrations would be unable to find a family of ancient primordial origin. However, these results still show promise for future work. The clustering statistic is still close to 1 radian, and the convergence of three family members could hint that further refinement of optimum integration parameters would yield a larger number of converging KBOs, potentially increasing the outer limit of family-finding backwards integrations to the age of the solar system.

Overall, the age constraints to backwards integration's family finding capabilities appear to not be a major hindrance. Indeed, it appears that backwards integration may be a feasible method of family detection for collisional families of more ancient origin than was originally expected. Integrations run with the current standard parameters should be able to successfully find families of ages on the order of 1 billion years or less, and further integration refinement could extend this boundary further into the past. Attempted rediscovery of Haumea family ($t_{age} = 3.5 \pm 2$ Gyr) via backwards integration would push this envelope, but theoretical young KBO families are well within the capabilities of discovery through backwards integration.

3.4 Exploring Numerical Chaos

A final hurdle associated with attaining accurate results from backwards integrations is that of numerical chaos. This is not a challenge unique to this research project, but plagues any computer calculation requiring high levels of precision, and can best be illustrated by resorting to the familiar analogy of the butterfly effect. In this scenario a the wing-flap of a small butterfly triggers a series of successively larger reactions ultimately resulting in a devastating large-scale typhoon.

Though initially created to describe difficulties in weather forecasting, it has direct application to long-term orbital integrations. Computer algorithms suffer from a common issue known as round-off error, in which numerical precision is lost by ascribing parameters to variable formats of specific bit-size. The resulting numbers rounded to fit the variable precision could vary from integration to integration, notwithstanding identical starting parameters. When dealing with integrations of lengths on the order of billions of years, these small round-off errors (i.e. numerical chaos) can snowball into massive orbital differences — even instabilities — in the long term.

Thus it was a real possibility that numerical chaos would prevent any integrations from achieving viability beyond certain timescales. To test the impact of chaos, two integrations were run that were largely identical to each other. However, in one integration the starting location of Uranus was shifted by 1 meter relative to the other integration. This small shift could be seen as the functional equivalent of the butterfly flapping its wings.

Analysis of these integrations then revealed them to be nearly identical, suggesting that numerical chaos is not a significant factor in long-term family integrations. This conclusion is bolstered by the ability of REBOUND to run bit-by-bit reproducible integrations, or in other words produce identical results when the same integration is run twice consecutively. Such an ability implies a sophisticated method of dealing with floating-point precision and round-off error, a feature which bodes well for minimizing the impact of numerical chaos.

Chapter 4

Conclusions and Discussion

This project began two years ago as a "reach-for-the-stars" endeavor. While backwards integration has yet to yield real-life results in the form of a new Kuiper belt family, its potential as a family-finding method is tantalizing.

4.1 Feasibility of Backwards Integration as a Family-finding Method

Based on the research and results presented in chapter 3, backwards integration seems to be a highly realistic means of family-finding for Kuiper belt families. Though backwards integrations in the asteroid belt are only viable through very short timescales (< 10 Myr), the relative dynamical stability of the Kuiper belt vastly increases the efficacy of backwards integration for KBOs.

The several challenges to backwards integration as a family-finding technique seem to not be formidable enough to preclude future success. Observational uncertainties, while currently large, are not prohibitively so, especially for timescales less than 1 Gyr. Assessing orbital elements in probabilistic terms rather than deterministic terms further reduces the impact of uncertainties, increasing the possible range of backwards integrations far beyond 1 Gyr. As new data from current and future telescope missions becomes available, observational uncertainty will only become less problematic.

Systematic uncertainties, while somewhat more problematic than their observational counterparts, are not the death knells they were once feared to be. Systematic uncertainty in integrations run with the current standard parameters appear to have an outer limit of 1 Gyr. Thus, younger KBO families with ages similar to the several known asteroid families should be easily found through backwards integration. Even families approaching primordial origin could feasibly discovered. Additionally, the 1 Gyr threshold imposed by systematic limitations could reasonably be extended by further investigation of the relative impact of various integration parameters or by running more representative integrations requiring more processing power.

Numerical chaos appears to not be a significant factor, and does not meaningfully limit integration timescales. Additionally, the bit-for-bit reproducability of REBOUND integrations is encouraging and could be of great utility to future research.

In sum, backwards integration appears definitively feasible. In a region populated with so many objects but about which so little is known in the way of familial composition, backwards integration could finally unlock the door to the dynamical state and evolution of the Kuiper belt.

4.2 Going Forward

As mentioned in section 4.1 and throughout chapter 3, further research devoted to refining and improving integration methodologies could possibly extended the stated thresholds further into the past. This effort represents the first track of future work: methodological refinement. While reasonable success has been attained with current integration parameters and techniques, additional hours spent dialing in the integrations could further reduce the impact of both observational and systematic uncertainties.

The second track of future work consists of formally commencing a search for KBO families using backwards integration techniques. Thus far all work has centered around testing and re-testing backwards integration as a technique, but few hours have been devoted to actually searching for families. The next obvious step then is to start searching in earnest.

In fact, this search has already begun. In early March 2019, a set of backwards integrations were run on the BYU Supercomputer Mary Lou, with a single integration created for every known Kuiper belt object. This set added to nearly 2000 total integrations, each run through 5 billion years. Included in each were the 4 Jovian planets, as well as 30 clones of the object in question.

Some analysis of these integrations has already begun. First steps have included extracting the orbital data for each timestep from the REBOUND SimulationArchives and storing them in Python .npy files. The mean values of each orbital element for each KBO at each timestep was then computed by averaging these values for every integrated clone (ejecting any that had entered unstable orbits). This has resulted in a massive amount of data (roughly 6.6 TB), which poses the significant challenge of determining how best to sift through it searching for families. The initial approach has involved creating two-dimensional arrays, comparing every KBO to every other KBO in Ω-space, and stacking these arrays in time. From these, a "close-ness" threshold could be chosen, and KBO pairs within that threshold could be identified. Subsequently, graph theory could be employed to examine these pairings and then identify clusters of pairings, from there using case-by-case analysis to identify families. This work is ongoing, and could possible produce ground-breaking results in the very near term.

The number of yet undiscovered KBO collisional families is unknown, though current research theorizes there are many (Barr & Schwamb 2016; Levison et al. 2008; Marcus et al. 2011). One can imagine dozens of these families quietly orbiting in the outer solar system, patiently awaiting the day they are finally detected. But regardless of the date on which the first of these many families is discovered, backwards integration figures to be a prominent part of the process. As the philosopher Kierkegaard once said: "The beginning is not that with which one begins, but at which one arrives at the beginning backwards."[1]

[1]Of course, the eminent Rudyard Kipling (after whom a crater on the planet Mercury is named) famously penned: "Never look backwards or you'll fall down the stairs." It is left as an exercise to the reader to determine the advice they follow.

Bibliography

Bannister, M. T., et al. 2016, AJ, 152, 70

Barr, A. C., & Schwamb, M. E. 2016, MNRAS, 460, 1542

Bottke, William F., J., Vokrouhlický, D., Rubincam, D. P., & Nesvorný, D. 2006, Annual Review of Earth and Planetary Sciences, 34, 157

Bottke, W. F., Vokrouhlický, D., Broz, M., Nesvorný, D., & Morbidelli, A. 2001, Science, 294, 1693

Bowell, E., Muinonen, K., Harris, A. W., & Howell, S. B. 1994, in AAS/Division for Planetary Sciences Meeting Abstracts, Vol. 26, AAS/Division for Planetary Sciences Meeting Abstracts #26, 27.11

Brown, M. E., Barkume, K. M., Ragozzine, D., & Schaller, E. L. 2007, Nature, 446, 294

Brown, M. E., Trujillo, C. A., & Rabinowitz, D. 2005, International Astronomical Union Circular, 8577, 1

Campo Bagatin, A., Benavidez, P. G., Ortiz, J. L., & Gil-Hutton, R. 2016, MNRAS, 461, 2060

Carruba, V., De Oliveira, E. R., Rodrigues, B., & Requena, I. 2018a, MNRAS, 479, 4815

Carruba, V., & Morbidelli, A. 2011, MNRAS, 412, 2040

Carruba, V., Vokrouhlický, D., Nesvorný, D., & Aljbaae, S. 2018b, MNRAS, 477, 1308

Celletti, A., Kotoulas, T., Voyatzis, G., & Hadjidemetriou, J. 2007, MNRAS, 378, 1153

Chiang, E. I. 2002, ApJ, 573, L65

Hirayama, K. 1918, AJ, 31, 185

Karlsson, O. 2011, Astronomische Nachrichten, 332, 562

Knezevic, Z., Lemaître, A., & Milani, A. 2002, The Determination of Asteroid Proper Elements, 603–612

Kondratyev, B. P. 2016, Ap&SS, 361, 169

Kuchner, M. J., Brown, M. E., & Holman, M. 2002, AJ, 124, 1221

Kuiper, G. P. 1974, Celestial Mechanics, 9, 321

Lacerda, P., Jewitt, D., & Peixinho, N. 2008, AJ, 135, 1749

Leinhardt, Z. M., Marcus, R. A., & Stewart, S. T. 2010, ApJ, 714, 1789

Levison, H. F., Morbidelli, A., Vokrouhlický, D., & Bottke, W. F. 2008, AJ, 136, 1079

Lockwood, A. C., Brown, M. E., & Stansberry, J. 2014, Earth Moon and Planets, 111, 127

Lykawka, P. S. 2012, Monographs on Environment, Earth and Planets, 1, 121

Marcus, R. A., Ragozzine, D., Murray-Clay, R. A., & Holman, M. J. 2011, ApJ, 733, 40

Nesvorný, D., Bottke, William F., J., Dones, L., & Levison, H. F. 2002, Nature, 417, 720

Nesvorný, D., & Bottke, W. F. 2004, Icarus, 170, 324

Nesvorný, D., Bottke, W. F., Levison, H. F., & Dones, L. 2003, ApJ, 591, 486

Nesvorný, D., Brož, M., & Carruba, V. 2015, Identification and Dynamical Properties of Asteroid Families, 297–321

Nesvorny, D., Vokrouhlicky, D., & Bottke, W. F. 2006, Science, 312, 1490

Ortiz, J. L., et al. 2012, MNRAS, 419, 2315

—. 2017, Nature, 550, 219

Proudfoot, B. C. N., & Ragozzine, D. 2019,

Rabinowitz, D., Schwamb, M. E., Hadjiyska, E., Rojo, P., & Tourtellottee, S. 2011, in EPSC-DPS Joint Meeting 2011, Vol. 2011, 1642

Rabinowitz, D. L., Barkume, K., Brown, M. E., Roe, H., Schwartz, M., Tourtellotte, S., & Trujillo, C. 2006, ApJ, 639, 1238

Ragozzine, D., & Brown, M. E. 2007, AJ, 134, 2160

—. 2009, AJ, 137, 4766

Rein, H., & Liu, S. F. 2012, A&A, 537, A128

Rein, H., & Tamayo, D. 2015, MNRAS, 452, 376

—. 2017, MNRAS, 467, 2377

Santos-Sanz, P., Ortiz, J. L., Aceituno, F. J., Brown, M. E., & Rabinowitz, D. 2005, International Astronomical Union Circular, 8577, 2

Schlichting, H. E., & Sari, R. 2009, ApJ, 700, 1242

Tsiganis, K., Knežević, Z., & Varvoglis, H. 2007, Icarus, 186, 484

Tsirvoulis, G. 2019, MNRAS, 482, 2612

Vokrouhlický, D., Brož, M., Morbidelli, A., Bottke, W. F., Nesvorný, D., Lazzaro, D., & Rivkin, A. S. 2006, Icarus, 182, 92

Vokrouhlický, D., et al. 2009, A&A, 507, 495

Volk, K., & Malhotra, R. 2012, Icarus, 221, 106

Winter, O. C., Borderes-Motta, G., & Ribeiro, T. 2019, MNRAS, 484, 3765

Wisdom, J., & Holman, M. 1991, AJ, 102, 1528

Zappala, V., Cellino, A., Farinella, P., & Knezevic, Z. 1990, AJ, 100, 2030

Zappala, V., Cellino, A., Farinella, P., & Milani, A. 1994, AJ, 107, 772