

The Strength of Smoothed Particle Hydrodynamics in Modelling Binary Interactions

By

Thomas Reichardt

Supervisors: Professor Orsola De Marco and Dr Jan Staff

A thesis submitted to Macquarie University
for the degree of Master of Research
Department of Physics and Astronomy
January 2016



MACQUARIE
University
SYDNEY • AUSTRALIA

Except where acknowledged in the customary manner, the material presented in this thesis is, to the best of my knowledge, original and has not been submitted in whole or part for a degree in any university.

Thomas Reichardt

Acknowledgements

First of all, I'd like to acknowledge my family. My parents and my brother have always supported me with my choice of career path, encouraged me to do my best in all of my ventures, and supported me for 22 years. I can't thank them enough for it. To Mum, thanks for pushing me when I start to procrastinate and telling me to give it my all. To Dad, thanks for patiently listening to me ramble on about my coding woes and introducing me to delicious coffee (which has become fundamental). Also, thanks Pete for being a great brother, and friend, and dealing with my abnormal affection.

Secondly, I'd like to acknowledge my supervisors, Orsola De Marco and Jan Staff. They have provided invaluable support, direction and insight throughout my entire project, helping me to grow as a researcher. Orsola has remained optimistic throughout the entire year about my progress, certainly beyond my own expectations, which provided an excellent working environment. She has always been there when I needed advice or answers; I can't think of any other who I would prefer as a supervisor or friend. Thanks also to Jan, who (until he left!) visited my desk every day to check if I had any questions that needed answering or interesting results that I wanted to show off, and patiently explained the points that probably should have been obvious to me. On a similar note, I'd like to acknowledge Roberto Iaconi, who introduced me to the code that I have been using and has, on more than one occasion, listened as I tossed ideas around. Finally, thanks to Daniel Price and James Wurster, who accommodated me twice at Monash University and, more importantly, wrote the code that I have been using this year for my work.

Lastly, thanks to all of my friends who have offered many distractions from the world of astrophysics. I am hesitant, but I will acknowledge Dan and Adrian (I suppose) for their roles of reminding me about deadlines and guidelines which I had not met. Thanks also to Tim and Mitch for the numerous evenings we've spent this year (and many years prior) hanging out; a welcome reprieve from working.

Abstract

The common envelope interaction gives rise to the formation of close binaries comprising at least one evolved star, such as a white dwarf. The idea behind the common envelope interaction is attractively simple. However, hidden complexities have prevented a full understanding of the interaction. Hydrodynamic simulations of this interaction have been pivotal in gaining an understanding of compact binary systems and phenomena that may result from them, such as novae and x-ray binaries. Unfortunately, notable disagreements still exist amongst simulations and between simulated parameters of post-common envelope binaries and observations, such as differing final separations.

In this work, we perform common envelope simulations with a new smoothed particle hydrodynamics code, PHANTOM, which has never been trained on this problem before. The lack of a simulation domain boundary, along with excellent conservation of energy and angular momentum, allows us to track the interaction in its entirety. We start by reproducing the simulations of Passy et al. (2012) to calibrate the code. PHANTOM is then used to carry out preliminary simulations of the phases preceding the fast inspiral and its effects on later stages of the interaction. We also investigated how bound gas that falls back onto the binary after the inspiral phase, leading to a new interaction. Although some computational issues remain to be resolved, PHANTOM has thus far proved to be an excellent tool in the study of binary interactions.

Contents

| | |
|---|------------|
| Acknowledgements | v |
| Abstract | vii |
| Contents | ix |
| 1 Introduction | 1 |
| 1.1 The formation of close, evolved binaries | 1 |
| 1.2 Theoretical description of the common envelope interaction | 3 |
| 1.3 Previous common envelope simulations | 6 |
| 1.4 Motivation for the current work | 8 |
| 2 SPH Simulations of the Common Envelope Interaction | 11 |
| 2.1 The smoothed particle hydrodynamics numerical method | 11 |
| 2.2 Lagrangian SPH and Eulerian grid codes | 14 |
| 2.3 The PHANTOM SPH code | 15 |
| 3 Benchmark Simulations | 19 |
| 3.1 Repeating the simulations of Passy et al. (2012) with PHANTOM | 19 |
| 3.1.1 Comparison of the evolution of the orbital separation | 22 |
| 3.1.2 The inspiral timescale | 23 |
| 3.1.3 Comparison of the energies and angular momenta | 24 |
| 3.2 Convergence tests | 28 |
| 4 The Pre-Common Envelope Phase | 33 |
| 4.1 The Roche-lobe overflow phase | 33 |
| 4.1.1 The corotating frame | 38 |
| 4.2 The pre-Roche lobe overflow phase | 42 |
| 5 Fallback of Gas During the Common Envelope Interaction | 45 |
| 5.1 Examination of fallback gas in the simulation with a $100 R_{\odot}$ initial separation | 45 |
| 5.2 Examination of fallback gas in the simulation with a $218 R_{\odot}$ initial separation | 48 |
| 6 Summary, Conclusions and Future Work | 51 |
| A Appendix A: Table of Simulations | 53 |
| B Appendix B: Movies and Code | 55 |
| References | 57 |

1

Introduction

1.1 The formation of close, evolved binaries

As telescopes and detectors have steadily improved and increasingly large swathes of sky have been surveyed over time, it has slowly become clear that the Universe is filled with transient events. Nova and supernova outbursts are some of the most well known and studied examples, due to the fact that they are frequent (the former) or bright enough to be visible at very large distances (the latter). Other types of transient events, such as x-ray binaries, mergers or other stellar exotica, are less frequent and still comparatively poorly studied, but data is rapidly accumulating (Kasliwal et al., 2012). Many of these events require close binary systems, and are expected to be formed from the result of a binary system undergoing at least one “common envelope” (CE) interaction phase. This type of interaction was originally proposed by Paczynski (1976) as a method of reducing the orbital separation of binary systems to explain how cataclysmic variables (stellar systems which irregularly experience dramatic increase in brightness before returning to their original state) may form.

Supernovae are a well known transient phenomenon, with the more relevant case here being type Ia supernovae (SNe Ia). These events have their origins in short period binary systems with either one or both of the components being a white dwarf. There are three common paths that have been suggested to lead to SNe Ia. The single degenerate path requires a white dwarf primary with a main sequence companion, with the two components being close enough that gas is being stripped from the outer atmosphere of the main sequence star and accreted by the white dwarf. Eventually, this will increase the mass of the white dwarf above the Chandrasekhar mass limit, at which point the pressure exerted due to self-gravity overcomes the electron degeneracy pressure in the white dwarf and the star tears itself apart in a thermonuclear explosion (Han and Podsiadlowski, 2004). The second scenario is known as the double degenerate scenario and occurs when two white dwarfs merge, pushing the total mass over the Chandrasekhar mass and detonating (Webbink, 1984). The final path is known as the double detonation scenario, wherein helium is accreted slowly and stably enough onto the surface of the white dwarf primary that a helium shell builds up around it. At some point, this shell will hit a critical mass and begin to fuse explosively,

sending a shock wave through the white dwarf which can cause the local pressure to overcome electron degeneracy pressure and cause a runaway explosion (Fink et al., 2007). The common trend here is that all of the scenarios require a highly evolved primary star that is very close to the companion.

SNe Ia are particularly useful as they are thought to have some standardisable peak luminosity, and hence can be used as distance indicators for far points in the Universe. Schmidt et al. (1998) and Riess et al. (1998) observed many of these supernovae at great distances, and applied their standardisable brightness to determine that the expansion of the Universe was accelerating. For this work, Brian Schmidt, Adam Reiss and Saul Perlmutter were awarded the 2011 Nobel Prize in Physics.

Novae and dwarf novae are two examples of a group of transient events known as cataclysmic variables. These variables are binary systems that experience occasional or semi-periodic, dramatic increases in brightness. Similarly to the SNe Ia progenitor systems, nova and dwarf nova outbursts result from systems wherein a white dwarf primary has formed an accretion disk from material stripped from a Roche lobe filling companion, where the Roche lobe is a teardrop shaped region of space around the star within which gas is gravitationally attracted to the star.

The close binaries from which these transient events arise are thought to be formed through a common envelope interaction. The more massive of the stars in an initial binary system will eventually evolve to have a radius on the order of magnitude of the separation between the two stars. Mass transfer leads then to the formation of a common envelope within which the secondary and the core of the primary will spiral closer together (Ivanova et al., 2013). The core of the giant star is essentially a white dwarf, so finding binary systems with at least one white dwarf wherein the orbital separation is shorter than the diameter of the red giant from which the white dwarf formed is a direct way of finding systems that must have undergone a common envelope event. Some of these systems are surrounded by a planetary nebula (Bond, 2000; Miszalski et al., 2009), the remnant of the common envelope that was shed by interactions with the cores.

Close binary systems where both components are either neutron stars or black holes are also thought to require one or more common envelope interactions in their formation. It is fairly simple to understand why the common envelope is thought to be a necessary part of the evolution of these types of systems. White dwarfs, neutron stars and black holes all require normal stellar evolution, where a star will progress from the main sequence to the red giant branch (RGB). Further stages of evolution may occur, but the RGB is the first stage from which a common envelope is a natural result. When moving from the main sequence to the RGB, a Sun-like star's radius will increase from $1 R_{\odot}$ to approximately $170 R_{\odot}$ (Sackmann et al., 1993), which is approximately two orders of magnitude larger than the orbital separations of observed post-common envelope binaries (Schreiber and Gänsicke, 2003; De Marco et al., 2011). So understanding the common envelope phase, and how the initial binary parameters lead to different outcomes, becomes necessary to be able to predict the formation rates of close binary classes, their parameters and transient phenomena in general.

The most commonly cited example, and probably the only example, that we have of a common envelope being observed in action is V1309 Scorpii. This object was discovered

in September 2008, though observations taken by the OGLE-III and OGLE-IV surveys fortuitously had covered the area of sky in which the object was situated for the prior 7 years. This system has been observed since 2001, with observations from this time revealing it to have been a close binary system with a period of approximately 1.44 days. This period decreased by about 0.02 days between the first observation and 2007 (Tylenda et al., 2011). The light curve during this period showed two distinct peaks, offset by roughly half a period, corresponding to two stars in a very tight, eclipsing orbit. In 2007 the secondary peak in the light curve quickly disappeared, indicating that the secondary star had been somehow obscured. In April of 2007, V1309 Sco had reached a local maximum in brightness, with an approximate magnitude of 15.5. At this point, the magnitude began to drop down to approximately 16.8 until, in March 2008, the brightness increased in a roughly exponential manner to a maximum magnitude of ~ 6.8 (Tylenda et al., 2011). Having these observations over a 7 year period allowed a complete picture of the orbital evolution of the system to be determined. After the massive outburst in 2008, there was no sign of binarity in the system, with the components presumably having merged to form one star. These observations leave little doubt that, prior to 2008, V1309 Sco was a contact binary. Throughout 2007, it seems that the system underwent a common envelope interaction, as evidenced by the fact that the second peak in the light curve disappeared. Finally, in 2008 the common envelope caused the two stars to spiral towards each other even more than previously, before finally merging and liberating gravitational energy in an eruption that increased the luminosity by 10^4 times. As there are no other mechanisms known to explain both the changing period of the system and the resulting huge outburst, simultaneously, the evolution of V1309 Sco was almost certainly caused by a common envelope event (Tylenda et al., 2011; Nandez et al., 2014).

A large fraction of all intermediate luminosity optical transients are today thought to be such mergers of two main sequence stars (Kasliwal, 2013; Soker and Kashi, 2011). If this is the case, then common envelope events are being readily observed even now. With increasingly sensitive observations in the years to come, these data will be obtained ever more frequently.

1.2 Theoretical description of the common envelope interaction

Although common envelope interactions are seldom observed, observational data of post-common envelope binaries offer a channel through which we can gain greater knowledge of the final parameters of such an interaction. There are numerous systems that are thought to have gone through a common envelope phase, as their current orbital separations are on the order of 1 to $10 R_{\odot}$, and the primary star is a white dwarf (De Marco et al., 2011; Schreiber and Gänsicke, 2003). Schreiber and Gänsicke (2003) show that almost half of a sample of 30 post-common envelope binaries will evolve to transfer mass again in the future. That is, within a Hubble time, magnetic braking and gravitational radiation will decrease their separations such that they will start transferring mass and evolve into cataclysmic variables.

The common envelope interaction began as an idea to explain the formation of cataclysmic variable systems, focusing particularly on the formation history of the short period binary V471 Tau (Paczynski, 1976). The common envelope phase is defined in that paper by first considering a system comprised of two main sequence stars with some relatively wide

separation. The more massive component of this system will eventually evolve to become a red giant and overflow its Roche lobe. When this occurs, the outer atmosphere of the giant streams to the companion and a contact binary system is formed. Paczynski (1976) goes on to suggest that material will likely build up around both stars, and that there is no foreseeable reason for the gas beyond the outer Lagrangian point to be shed without interactions with the two cores. This envelope may last for many orbital periods, and is known as a common envelope. Interactions of the stars with this envelope dictate a decrease in orbital separation, and a lifting of the envelope material. Paczynski (1976) also mentions private conversations with Jeremiah Ostriker as well as Ron Webbink's PhD thesis as the genesis of the idea.

Descriptions of the common envelope are often simplified and broken down into several distinct phases. Chronologically, the first phase is essentially the evolution of the system that leads up to the common envelope event. Primarily, it explains how a corotating, stable binary may be transformed into an inspiralling binary. This stage incorporates the stellar evolution of the primary. At some point, the star will cease fusing hydrogen in the core and begin to contract slightly. This contraction allows the ignition of a hydrogen fusing shell around the helium core, the result of which is that the star will evolve into a red giant, increasing its radius until it is approximately 100 times larger than before. It is likely that this expansion will cause the outer atmosphere of the primary to expand beyond its Roche lobe, setting up a period of mass transfer and, if the donor star is more massive than the companion, the system experiences a resulting decrease in the orbital separation until the next stage can occur. Another potential route by which a corotating binary may decrease its orbital separation is known as the Darwin instability (Darwin, 1879). In the case of a circular orbit, when the spin angular momentum of the stars is greater than one-third of the orbital angular momentum, or:

$$J_{spin} > \frac{1}{3} J_{orb} \quad (1.1)$$

and the orbital separation is reduced, angular momentum is removed from the orbit. The reduction in the orbital separation increases the angular velocity, and the result is that the lost orbital angular momentum is added to the spin angular momentum of the stars in an attempt to keep the primary spin synchronised with the orbit. However, when the above threshold criterion is met, keeping the binary in corotation requires more angular momentum than what was extracted by the shortening of the orbit, so the system again reduces its separation to accommodate this. This causes a further need to spin up the primary, resulting in a positive feedback loop which quickly extracts angular momentum from the orbit, forcing the two stars closer together.

The second stage of the common envelope is often known as the fast inspiral. This stage is characterised by the secondary being enveloped by the atmosphere of the primary, resulting in a sudden and dramatic reduction in the orbital separation. However, the reduction of the orbital separation means that the envelope has gained some angular momentum, as the gas interacts gravitationally with the cores and is essentially lifted from the system. It is the interaction between the cores and the gas that drives the quick evolution of this system.

The plunge-in of the secondary star into the envelope of the primary will likely eventually level out into a much slower, self-regulated inspiral phase, wherein the frictional luminosity released by the inspiral pushes gas away, which reduces the luminosity and hence the local density increases again (Meyer and Meyer-Hofmeister, 1979). This sequence regulates the

further inspiral of the cores, and operates on the thermal timescale of the system (Ivanova et al., 2013).

At this point, it is important to note that there are two possible outcomes of a common envelope interaction. Either the two stars will end up in a much more compact orbital arrangement than when they started and the envelope will be completely removed from the system, or else, if the envelope is not ejected, the cores continue to transfer angular momentum to the envelope and continue to inspiral until they end up merging. It is this second case that was the final outcome of V1309 Sco. Despite the apparent simplicity of this interaction, we are still unclear on how the final outcomes (e.g. final separation) may depend on the initial parameters.

A key to understanding common envelope evolution is the nature of the interaction between the cores (i.e. the companion star and the degenerate helium or carbon-oxygen core of the primary star) and the common envelope material. At first, it seems plausible that the angular momentum of the cores is reduced as a result of direct collisions of the primary core and companion with the surrounding gas. While this effect seems to be a reasonable explanation in principle, in practice, the cross-sections of the cores are nowhere near large enough for collisions alone to reproduce the physical parameters of post-common envelope objects. This type of interaction, known as frictional or “hydrodynamic” drag, is certainly present in the common envelope, however, simulations have found that it is dwarfed by a second drag force known as “gravitational” drag (Ricker and Taam, 2008). Indeed, Passy et al. (2012) found with their simulations of the common envelope that the hydrodynamic drag must be around three orders of magnitude smaller than the gravitational drag. Thus, when dealing with the common envelope, the hydrodynamic drag is quite often ignored.

As suggested by the name, gravitational drag does not require any physical contact between two objects, instead dictating a transfer of energy and angular momentum purely through the influence of gravity. The exact nature of gravitational drag in the presence of a density gradient has never been fully studied, though MacLeod and Ramirez-Ruiz (2015) briefly addressed it in their study of asymmetric accretion flows. As the core and companion star move through the gas of the envelope, the gas flows around them, forming a dense wake behind them. This contributes to slowing the cores down and to reducing the orbital separation as a result. On the other hand, in the common envelope simulation of Staff et al. (2015), the two cores maintain an approximately Keplerian velocity throughout the inspiral. It is thought, therefore, that a more important effect in this case is that the density is slightly higher towards the centre of the giant star’s envelope, and hence more gas will move past the inspiralling companion on the inside. This causes an asymmetry of gas flowing around the star. This imbalance causes a slight shift in the trajectory of the cores, forcing them into a tighter orbit. The gas gains the angular momentum lost from the orbit, and is typically lifted from the system. Iben and Livio (1993) define a formula with which to estimate the magnitude of the gravitational drag:

$$F_{\text{drag}} = \xi \pi R_a^2 \rho (v - v_e)^3 \quad (1.2)$$

where ξ is a scaling factor that depends on the Mach number (see also Ostriker (1999)), ρ is the density of the gas, $v - v_e$ is the relative velocity between the cores and the gas and R_a is

the accretion radius defined by:

$$R_a = \frac{2GM}{(v - v_e)^2 + c_s^2},$$

where c_s is the local sound speed.

The common envelope interaction has generally been described with the ‘energy formalism’, which is written in the simplest form by equating the gravitational binding energy of the envelope with the difference in orbital energy of the initial and final configurations of the two stars:

$$E_{\text{bind}} = \Delta E_{\text{orb}} \equiv E_{\text{orb,i}} - E_{\text{orb,f}}. \quad (1.3)$$

In common envelope descriptions, an envelope ejection efficiency parameter, α , is also inserted, dictating the efficiency with which the orbital energy unbinds the envelope. The common definitions of binding energy and the change in orbital energy are:

$$E_{\text{bind}} = G \frac{M_1 M_{1,\text{env}}}{\lambda R_1} \quad (1.4)$$

$$\Delta E_{\text{orb}} = -G \frac{M_1 M_2}{2a_i} + G \frac{M_{1,\text{c}} M_2}{2a_f} \quad (1.5)$$

where G is the gravitational constant, M_1 is the mass of the primary, $M_{1,\text{env}}$ is the envelope mass, $M_{1,\text{c}}$ is the primary core mass, M_2 is the companion mass, R_1 is the radius of the primary, a_i and a_f are the initial and final orbital separations and λ is a parameter characterising the structure of the primary star with values generally between 0.2 and 0.8 for main sequence stars (Dewi and Tauris, 2000). The value of λ describes the distribution of the mass within the star, with lower values characterising more compact stars and larger values extended stars. Hence, the common form of the energy formalism is (Ivanova et al., 2013):

$$G \frac{M_1 M_{1,\text{env}}}{\lambda R_1} = \alpha \left(-G \frac{M_1 M_2}{2a_i} + G \frac{M_{1,\text{c}} M_2}{2a_f} \right). \quad (1.6)$$

Other forms of this equation have been used in studies of the common envelope (De Marco et al., 2011), which inevitably leads to confusion in comparison of results.

1.3 Previous common envelope simulations

Simulations of the common envelope interaction are an important tool with which we can understand observed systems. They have been carried out over many years, but have previously suffered from being under-resolved. While this problem still persists, simulations are quickly increasing in resolution. One of the first notable simulations of the common envelope phase was carried out by Rasio and Livio (1996). This simulation was conducted with a smoothed particle hydrodynamics (SPH) code with approximately 5×10^4 particles, and modelled the dynamical mass transfer and fast inspiral stages of a common envelope interaction. Their binary system contained a $4 M_\odot$, $63 R_\odot$ red giant primary and a $0.7 M_\odot$ point-mass companion, separated by $100 R_\odot$. Approximately 10% of the mass of their envelope was unbound from the binary, though they noted that calculating the bound mass with a term for the enthalpy of each particle returns a slightly larger fraction of unbound mass than

| Simulation | M_1 (M_\odot) | M_2 (M_\odot) | R_1 (R_\odot) | a_i (R_\odot) | a_f (R_\odot) | M_{unb} (M_\odot) |
|-------------------------|---------------------|---------------------|---------------------|---------------------|---------------------|--------------------------------|
| Rasio and Livio (1996) | 4 | 0.7 | 63 | 100 | 2 | 8-14% |
| Sandquist et al. (1998) | 3, 5 | 0.4, 0.6 | 200, 356 | 289, 536 | 4, 9 | 31%, 23% |
| Ricker and Taam (2012) | 1.05 | 0.6 | 32 | 61 | 9 | 26% |
| Passy et al. (2012) | 0.88 | 0.6 | ~ 90 | ~ 90 | 20 | 15% |
| Nandez et al. (2014) | 1.5 | 0.16 | ~ 3.3 | ~ 6.3 | 0 | $\sim 2\%$ |

Table 1.1: Parameters of the simulations carried out in previous works.

when enthalpy is ignored. The low fraction of unbound mass is one of the primary disagreements between simulations and observations. It is clear that the majority of the envelope mass must be unbound from the system, or else common envelope interactions would leave behind no close binary systems, only mergers.

Every simulation since the preliminary work of Rasio and Livio (1996) has encountered a similar problem. One of the simulations of Sandquist et al. (1998) consisted of a $3 M_\odot$, $200 R_\odot$ radius primary and a $0.4 M_\odot$ companion separated by $\sim 290 R_\odot$. The core of the primary was modelled as a point particle of mass $0.7 M_\odot$. These simulations, which were carried out with a nested Eulerian grid code, resulted in unbound masses in the range of 20-30% (note that this is an inference from their plots). Further simulations using SPH (Passy et al., 2012, to be discussed further in subsequent chapters) yielded an unbound mass percentage of 15%. Ricker and Taam (2012) ran simulations using the adaptive mesh refinement (AMR) grid code, FLASH (Fryxell et al., 2000), with a primary star of $\sim 30 R_\odot$ and mass $1.05 M_\odot$ and a $0.6 M_\odot$ companion at a separation of $\sim 60 R_\odot$. The core of the primary ($0.39 M_\odot$) and the companion were modelled with clouds of 200 000 particles moving as solid bodies, and interacting with the gas only gravitationally. This simulation yielded a similar unbound mass percentage compared to previous simulations (26%). Some more recent simulations have been performed by Nandez et al. (2014) using an SPH code, with the explicit aim of attempting to recreate the V1309 Sco merger. They used a binary system with a subgiant primary of mass $1.52 M_\odot$ and radius $3.3 R_\odot$, and a compact companion of mass $0.16 M_\odot$, separated by $6.3 R_\odot$. This simulation resulted in a merger and the unbound mass percentage was very low (2%). Some of the parameters of these simulations can be seen clearly in Table 1.1.

Some solutions to the problem of too little unbound mass have been proposed. A primary example is that of the role of the recombination energy in the energy budget for the common envelope interaction, which has most notably been addressed by Ivanova et al. (2015) and Nandez et al. (2015). As much of a giant star comprises ionised hydrogen and helium, there is a lot of stored energy that may be released when the plasma recombines to form neutral hydrogen and helium. Webbink (2007) calculated that the recombination energy of a giant star would be on a similar order of magnitude to the gravitational binding energy, and hence could potentially influence the common envelope interaction. Nandez et al. (2015) gives details of SPH simulations that include a recombination term in the energy formalism. For the most part, these simulations used approximately 10^5 particles. It was found that, by the end of the simulations, 99.8% of the envelope was unbound, and the orbital separations were $\sim 2 R_\odot$. This is a particularly promising result for the field, though the work is clearly still preliminary, in that the system they simulated was somewhat more compact than those simulated by other authors, so it is not possible to compare results. Further, their simulations were

mostly carried out with $\sim 100\,000$ particles (one simulation had $200\,000$ particles), which is a low resolution, and a convergence test was not carried out.

What is interesting about the results of Nandez et al. (2015) is that the final orbital separation is low. For the mass ratio $q \equiv \frac{M_2}{M_1} = 0.37$, their orbital separation decreased from an initial value of $60\,R_\odot$ to the final separation of $2\,R_\odot$. All past simulations, except for one, have separations that are large compared to the orbital separations of post-common envelope binaries (see Fig. 17 in Passy et al. (2012)) and are correlated with q . The only simulation prior to that of Nandez et al. (2015) to have a separation that was very low for the value of q is that of Rasio and Livio (1996), who had a final separation of $a_f = 2\,R_\odot$ with $q = 0.18$. It is not clear why some simulations have such low separations, nor is it clear whether they would observe the dependency on q seen by Passy et al. (2012), which is not found in observations. It is possible that low final separations are favoured in interactions with the heavy and/or compact giants used both by Rasio and Livio (1996) and by Nandez et al. (2015). What seems clear is that the addition of the recombination energy by Nandez et al. (2015), which solves the bound envelope problem, is not likely to be the reason for a low final separation; the addition of recombination energy to unbind more mass should, if anything, contribute to a larger, not smaller, final separation.

1.4 Motivation for the current work

Despite the importance of the common envelope interaction and the large body of work that has been dedicated to understanding it, there is still so much uncertainty in our knowledge of how common envelope interactions start, progress into the dynamical infall and terminate, leaving behind the close binary system.

Up until this point in time, there have been few simulations investigating what may happen before the rapid infall phase, during tidal interaction and Roche-lobe overflow phases, or what happens to gas that is ejected from the system during the rapid infall but remains bound. Kashi and Soker (2011) and De Marco et al. (2011) suggest that this bound material will eventually return to the system to form a circumbinary disc. Part of this disc may interact again with the binary system, and provide a new opportunity for the stars to transfer orbital angular momentum to the gas. This might have the effect of increasing the radius of the disc, while also decreasing the orbital separation of the binary (Kashi and Soker, 2011; Kuruwita et al., 2015). The post-fast inspiral phase of the common envelope interaction is particularly poorly simulated as it occurs over longer timescales than the fast inspiral phase. It may be particularly useful to explain why simulated final separations are larger than observations. Kuruwita et al. (2015) showed that the orbital separation does decrease during an interaction with a fallback disc, though the fallback interaction results in minimal further unbinding. The simulations of Nandez et al. (2015), on the other hand, provide us with a likely explanation of how more mass can become unbound. However, they do not readily explain how to reduce the post-rapid infall separation. Their simulations have low final separations, but this may be due to their using relatively compact primary stars and low initial orbital separations. It is possible that the actions of recombination energy and fallback gas in tandem may provide a better model for the common envelope interaction.

The simulations of Kuruwita et al. (2015) were carried out with a grid code and the

fallback setup was somewhat artificial. An SPH code is needed to be able to simulate the full common envelope interaction, from the lifting of the envelope to the formation of a circumbinary disc due to fallback gas, because grid codes will lose information on any matter that moves beyond the simulation boundary. For this reason, SPH codes are also better at tracking energy and angular momentum losses that inevitably plague numerical simulations. In general, SPH codes will perform comparably to grid codes in the common envelope parameter space, as confirmed by Passy et al. (2012); ultimately, both code types have advantages and disadvantages and it remains an asset to use both in tandem. A more detailed comparison of the two code types is offered in Section 2.2.

In this work, we will be using a new SPH code, PHANTOM, developed by Price et al. (2015), for our simulations. This was chosen because the code includes all the necessary physics and setup options, and because any modifications can be easily implemented with the help of the primary code developers. We have worked with Dr Price throughout this project in order to improve the ability of the code to simulate common envelope interactions. PHANTOM has also been developed specifically with non-cosmological applications in mind (Price et al., 2015), and while there are other SPH codes that have already been set up for common envelope interactions, training new codes at existing problems can help to shed new light on the matter.

One strong motivation for our work and for intensifying efforts on modelling the common envelope interaction is that time-domain surveys are becoming increasingly available and they should be able to detect and observe common envelope events as they occur. Without a reasonably accurate model of common envelope interactions, the data produced by such surveys will be difficult to interpret.

The Palomar Transient Factory was one of the early time-domain surveys, specifically built with the ability to investigate optical transient and variable events with a cadence varying from minutes to years (Rau et al., 2009). The design goals of the survey was specifically to observe a wide variety of transient events, some of which are the end results of common envelope interactions. Originally intended to operate from 2009 to 2013, a new collaboration was formed in order to keep the survey operating, until at least 2017 when its successor, the Zwicky Transient Facility, is scheduled to open (Surace et al., 2015). The Catalina Real-Time Transient Survey is a survey carried out by a collaboration between three wide-field telescopes in America and Australia. The survey covers an area of 30,000 square degrees with cadences of 10 minutes and greater. It has been particularly successful in identifying supernovae and cataclysmic variables, and all data is published immediately electronically and is freely available (Djorgovski et al., 2011). Finally, the Large Synoptic Survey Telescope is an 8.4 metre telescope that will undertake a 10 year survey of 30,000 square degrees of the sky, covering this area every three nights. The cadence of the imaging will vary such that objects changing in brightness on timescales ranging from minutes to decades will be sampled appropriately (Ivezic et al., 2008). This survey will offer an unparalleled combination of depth and sheer amount of data on transient events.

In addition to transient surveys, the advanced Laser Interferometer Gravitational Wave Observatory is slated to open in 2015, and will perhaps detect gravitational waves for the first time (Harry et al., 2010). Aside from the obvious importance of detecting gravitational waves as a test of General Relativity, this could provide data on X-ray binaries and other close

binary mergers comprising neutron stars and black holes. These are all excellent sources of gravitational waves when they merge, and all form via one or two common envelope interactions. In order to predict merger rates and hence gravitational wave source rates, population synthesis codes need to incorporate the efficiency of the common envelope ejection, a number (or function) that can potentially be determined via simulations.

SPH Simulations of the Common Envelope Interaction

In the previous chapter, we discussed the background of the common envelope interaction and explained why it is a worthwhile phenomenon to study. In particular, we outlined some of the many simulations that have been carried out in the past in this area. In this chapter, we will give some general details of the smoothed particle hydrodynamics method, compare the Lagrangian SPH and Eulerian grid methods and also examine the code, PHANTOM, used in this work.

2.1 The smoothed particle hydrodynamics numerical method

Smoothed particle hydrodynamics (SPH) is one of a couple of major classes of computational hydrodynamics techniques. It is a method of simulating fluid dynamics by solving the equations of motion for discrete parcels of the fluid, and hence it fits into a Lagrangian methodology of fluid flow. The primary difference between Lagrangian and Eulerian representations of the motions of fluid is that the Eulerian representation follows how the fluid moves through a fixed volume in space, while the Lagrangian method follows how a parcel of the fluid moves with respect to other parcels. Primarily, SPH provides an answer to the issue of computing the density from a collection of particles (Price, 2011). Each of the particles in the simulation has a constant mass, but has a volume that is inversely proportional to the number of particles in its immediate area. It is this latter property that is a unique function of SPH codes, and is what leads to the ‘smoothed’ nature of the particles. This means that each of the particles has their mass smoothed over a volume in a manner dictated by a smoothing *kernel*, W . Hence the density at any point can be written as a sum of the masses in the area weighted by the smoothing kernel (Monaghan, 2005; Price, 2011),

$$\rho(\mathbf{r}) = \sum_j^{N_{\text{neigh}}} m_j W(|\mathbf{r} - \mathbf{r}_j|, h) \quad (2.1)$$

where m_j , \mathbf{r}_j and h are the mass, position and smoothing length of particle j , respectively. The smoothing length of a particle dictates how spread out the particle is, such that low density particles will have a large smoothing length. Particles, then, are not point masses, but rather interpolation points in a continuous density field. The summation is carried out over the particles in the area surrounding the position at which the density is to be calculated; hence N_{neigh} is the number of neighbouring particles, and is related to the smoothing length and the chosen kernel. In the case of the most common kernel, the cubic spline (discussed below), the optimal number of neighbours is somewhere around 60, as particles begin to become paired when $N_{\text{neigh}} > 65$ (Price, 2011). More generally, a quantity which is a function of position in space, designated here as $A(\mathbf{r})$, can then be approximated at any point in the computational domain by solving the summation of this quantity over nearby mass elements,

$$A(\mathbf{r}) = \sum_{j=1}^{N_{\text{neigh}}} m_j \frac{A_j}{\rho_j} W(|\mathbf{r} - \mathbf{r}_j|, h) \quad (2.2)$$

It can be seen that if we are trying to calculate the density, $A_j = \rho_j$ and the sum becomes equation 2.1. This sum should be over all particles in the simulation, but the smoothing kernel is often chosen to have “compact support”. In other words, W decreases to zero beyond a radius of twice the smoothing length h from particle j as the kernel decreases enough to be negligible beyond $2h$, and the summation over all particles would drastically increase the time required to complete the simulation. Hence, the sum typically only includes the particles within $2h$ of the central particle.

The smoothing kernel function must obey a number of properties for it to be considered viable. The most general properties include (Liu et al., 2003):

1. The function must be normalised, that is, the integral of the function must be equal to unity;

$$\int W(|\mathbf{r} - \mathbf{r}'|, h) d\mathbf{r}' = 1$$

2. The function should be monotonically decreasing, and
3. The function should display the behaviour of the Dirac delta function in the limit that $h \rightarrow 0$;

$$\lim_{h \rightarrow 0} W(|\mathbf{r} - \mathbf{r}'|, h) = \delta(\mathbf{r} - \mathbf{r}')$$

While these conditions allow the possibility of the smoothing kernel being represented by an infinite number of functions, the most commonly used is the cubic spline, which takes the form

$$W(r, h) = \frac{1}{4\pi h^3} \begin{cases} (2 - q)^3 - 4(1 - q)^3 & \text{for } 0 \leq q \leq 1 \\ (2 - q)^3 & \text{for } 1 \leq q \leq 2 \\ 0 & \text{for } q > 2 \end{cases} \quad (2.3)$$

where $q \equiv \frac{r}{h}$. It should be noted that here we have written $\mathbf{r} - \mathbf{r}' = r$ as the kernel should be independent of direction, instead depending only on radius. This kernel is defined by $W(r, h) = \frac{1}{h^3} M_4$, where M_4 is from the group of M_n splines, identified by the Fourier transform (Monaghan, 2005)

$$M_n(x, h) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \left(\frac{\sin(kh/2)}{kh/2} \right)^n \cos(kx) dk$$

With the use of equations 2.1 and 2.3, the density may be calculated at any particle, and indeed at any point in the computational domain. With this knowledge, we can move on to address the question of what to do with it. The hydrodynamic equations dictate the properties of a fluid, and density is a typical input. The first of the fluid equations is the continuity equation, which enforces conservation of mass within a system:

$$\frac{d\rho}{dt} + \rho \nabla \cdot \mathbf{v} = 0, \quad (2.4)$$

where ρ is the density and \mathbf{v} is the velocity and the derivative here is known as the Lagrangian derivative, defined as:

$$\frac{da}{dt} = \frac{\partial a}{\partial t} + \mathbf{v} \cdot \nabla a,$$

for some variable a . The equation of motion of the gas is the next equation to be considered, and can be written in the following form:

$$\frac{d\mathbf{v}}{dt} = -\frac{1}{\rho} \nabla p - \nabla \Phi, \quad (2.5)$$

where p is the pressure, and Φ is the gravitational potential (hence, $\nabla \Phi = -\mathbf{g}$, for \mathbf{g} the gravitational acceleration). This equation represents conservation of momentum within the system. The third conservation equation expresses conservation of energy, and may be written:

$$\frac{du}{dt} = -\frac{1}{\rho} \nabla \cdot (p\mathbf{v}), \quad (2.6)$$

where u is the specific internal energy of the gas. The final equation, to complete the set, is the equation of state for the gas:

$$u = \frac{1}{\gamma - 1} \frac{p}{\rho}, \quad (2.7)$$

which gives a relationship between the specific internal energy and the pressure and density. In this equation of state, γ is the ratio of the heat capacities at constant pressure and constant volume. The most common value of γ is $\frac{5}{3}$ for an ideal, monatomic gas. The inclusion of radiation to the mix decreases the value of γ to approach the isothermal value of $\frac{4}{3}$. This equation of state ceases to be valid when an ideal gas is no longer being simulated; for example, the recombination of ions and electrons releases energy within the gas, changing the equation of state to a non-analytical form.

Equations 2.4 to 2.7 are the basis of fluid dynamics, and are discretised for use in a hydrodynamic code. In SPH, the continuity equation can be represented as a sum over the particles in the simulation in the following manner (Monaghan, 2005):

$$\frac{d\rho_a}{dt} = \rho_a \sum_b \frac{m_b}{\rho_b} \mathbf{v}_{ab} \cdot \nabla_a W_{ab}, \quad (2.8)$$

where ρ_a is the density at particle a , m_b and ρ_b are the mass and density of a neighbouring particle b , where the sum is carried out over all neighbouring particles. Here we have also written $\mathbf{v}_{ab} = \mathbf{v}_a - \mathbf{v}_b$ for the difference in velocities between particle a and neighbouring particle b , and $W_{ab} = W(\mathbf{r}_a - \mathbf{r}_b, h)$ is the value of the kernel of particle b at point a . The

SPH forms of the acceleration and energy equations can also be shown to be of the following forms:

$$\frac{d\mathbf{v}_a}{dt} = - \sum_b m_b \left(\frac{P_b}{\rho_b^2} + \frac{P_a}{\rho_a^2} \right) \nabla_a W_{ab} \quad (2.9)$$

$$\frac{du_a}{dt} = \frac{P_a}{\rho_a^2} \sum_b m_b \mathbf{v}_{ab} \cdot \nabla_a W_{ab} \quad (2.10)$$

Note that these equations in their raw form do not contain a term for dealing with gravity. This can be added, at which point these equations will exactly mirror their continuous counterparts.

SPH in this form is non-dissipative, meaning that there is no dissipation or diffusion of energy. While in general, this can be a good thing as energy is explicitly conserved in the system, some schemes require some form of artificial viscosity which must be added to the code (Price, 2011). This artificial viscosity often takes the form of an extra parameter added into the acceleration equation:

$$\frac{d\mathbf{v}_a}{dt} = - \sum_b m_b \left(\frac{P_b}{\rho_b^2} + \frac{P_a}{\rho_a^2} + \Pi_{ab} \right) \nabla_a W_{ab}, \quad (2.11)$$

where the viscosity term is defined by (Morris and Monaghan, 1997):

$$\Pi_{ab} = \begin{cases} -\frac{\alpha h \mathbf{v}_{ab} \cdot \mathbf{r}_{ab}}{\bar{\rho}_{ab} |\mathbf{r}_{ab}|^2} \left(\bar{c}_{ab} - 2 \frac{h \mathbf{v}_{ab} \cdot \mathbf{r}_{ab}}{|\mathbf{r}_{ab}|^2} \right) & , \text{ for } \mathbf{v}_{ab} \cdot \mathbf{r}_{ab} < 0 \\ 0 & , \text{ otherwise} \end{cases} \quad (2.12)$$

Here, α is a unitless viscosity parameter, c is the sound speed, with \bar{c}_{ab} denoting an average of c at positions a and b . To be able to describe shocks correctly, α must be relatively large. However, this can cause unwanted and unphysical dissipation in other areas of the simulation. Hence, viscosity switches have been devised to be able to artificially increase the viscosity in regions that require it. The standard prescription for this type of switch was devised by Morris and Monaghan (1997) and dictates the evolution of α with the equation

$$\frac{d\alpha}{dt} = \max(-\nabla \cdot \mathbf{v}, 0) - \frac{\alpha - \alpha_{\min}}{\tau} \quad (2.13)$$

where $\alpha_{\min} = 0.1$ is the generally assumed value, and τ is the e-folding time for α to reduce to the minimum value. This switch ensures that shocks are captured considerably better than with a static α , while also offering low dissipation in other regions.

2.2 Lagrangian SPH and Eulerian grid codes

Grid based codes have been used for common envelope simulations quite successfully. As mentioned in Section 2.1, they employ the Eulerian view of hydrodynamics, which models a fluid by calculating its bulk movement through stationary cells in space. These cells are arranged along a grid, and have constant volumes and contain variable masses. This is directly in opposition to an SPH code, which uses the Lagrangian perspective of modelling the

movement of fluid parcels, such that SPH particles have variable volume and constant mass.

As a result of grid codes having cells with fixed volume and position, there is a defined boundary to the simulation. If anything in the simulation moves beyond this boundary, then all information about it is lost. As may be expected, this can cause issues in the resulting morphology of a common envelope interaction, as well as making it somewhat difficult to calculate the fraction of bound mass or checking energy and angular momentum conservation. Further, when simulating a star, grid codes necessarily need to include a low density, high temperature medium surrounding the star in order to keep it stable. This can make it difficult to interpret some of the results of the simulation. These issues are easily solved in SPH, as the movement of the gas relative to an arbitrary origin is tracked by moving interpolation points, hence there is no set boundary and mass will never be lost. Further, as will be discussed in Chapter 3, the stabilisation of a star in SPH does not require an external medium to keep the star in equilibrium.

SPH codes, on the other hand, have drawbacks not present in Eulerian grid codes. For example, SPH codes will generally under-resolve the low density portions of a simulation; in the case of a common envelope simulation, this is initially the outside of the star and later the material that is ejected. This is because particles have their smoothing lengths grow in low number density areas, hence a large volume of space can be covered by a single particle and any underlying structure is not resolved. Further, the artificial viscosity of SPH may lead to dissipation issues in the simulations. Grid codes have a set resolution throughout the computational domain in unigrid, hence features of all densities are captured at a standard resolution. In adaptive mesh refinement grid codes, the resolution can be improved in areas of interest, where it is possible to specify the conditions for refinement. As we have explained in Section 1.4, the complementary advantages and comparable results of grid and SPH codes promote the use of both code types.

2.3 The PHANTOM SPH code

PHANTOM is a SPH code developed specifically with the design goals of being parallelised, modular and requiring a low memory footprint (Price et al., 2015). It has been developed over several years to be used in a variety of astrophysical applications. Modules involving magnetohydrodynamics, self gravity, chemistry and other facets useful over a range of astrophysical and hydrodynamic projects are readily available for implementation. Magnetic fields and H_2 chemistry are not considered in this project, and hence will not be covered here, however, self-gravity is integral for common envelope evolution, as much of the mass is tied up in non-static SPH particles.

PHANTOM writes the discretised, SPH form of the fluid acceleration equation as:

$$\frac{d\mathbf{v}_a}{dt} = - \sum_b m_b \left[\frac{P_a}{\rho_a^2 \Omega_a} \nabla W_{ab}(h_a) + \frac{P_b}{\rho_b^2 \Omega_b} \nabla W_{ab}(h_b) + \Pi_{\text{shock}}^a \right] + \mathbf{a}_{\text{ext}}(\mathbf{x}_a, t) + \mathbf{a}_{\text{sink-gas}}^a + \mathbf{a}_{\text{selfgrav}}^a, \quad (2.14)$$

where the final three terms give contributions to the acceleration owing to external forces, interactions between sink particles (to be described below) and gas, and self gravity, respectively. The term Ω_a is also a common addition to the acceleration equation, and is related to

the gradient of particle smoothing lengths:

$$\Omega_a \equiv 1 - \frac{\partial h_a}{\partial \rho_a} \sum_b m_b \frac{\partial W_{ab}(h_a)}{\partial h_a},$$

where the term $\frac{\partial h_a}{\partial \rho_a} = -\frac{3h_a}{\rho_a}$. PHANTOM also employs a slightly different version of the energy equation than what is described above in Equation 2.10:

$$\frac{du_a}{dt} = \frac{P_a}{\rho_a^2 \Omega_a} \sum_b m_b \mathbf{v}_{ab} \cdot \nabla_a W_{ab}(h_a) + \Lambda_{\text{shock}} + \Lambda_{\text{cool}} \quad (2.15)$$

where Λ_{shock} and Λ_{cool} are terms that are added to account for energy dissipation at the shock fronts, and for cooling of the gas, respectively (Price et al., 2015).

Like some other SPH codes, PHANTOM employs a *kd*-tree for finding neighbouring particles. The implementation in PHANTOM is very similar to the method described in Gafton and Rosswog (2011). The basis of this *kd*-tree is that the all of the particles in the simulation are grouped in a cell known as the ‘root’ cell. After this grouping, the root cell is split through the centre of mass with a plane perpendicular to the longest axis. This approximately leaves equal numbers of particles in each subsequent cell. This process is repeated recursively, splitting cells through their centre of mass until a particular cell has less than an arbitrary limit of N_{min} particles. These cells are known as ‘leaf nodes’. Neighbouring particles are then found by determining nodes that may possibly interact with each other.

Self-gravity in PHANTOM is approached in the standard fashion of having short range interactions computed with a direct summation of the kernel softened gravitational interaction between a target particle and its neighbours. Longer range interactions are introduced by addition of an acceleration between the centre of two leaf nodes, taking into account the quadrupole moment of the node. The long-range gravitational acceleration on particles in the leaf node is then calculated by using a second-order Taylor series expansion of the acceleration on the node, about the node centre.

Sink particles are a feature of PHANTOM important for our common envelope simulations. As will be described in Chapter 3, they are used to approximate the core of an RGB star as well as the compact companion. Typically, sink particles in SPH codes will be used to accrete the gas around them when certain conditions are met, however, in the case of our simulations, the accretion radius has been set to 0. Thus, although we refer to the cores as sink particles, they actually behave like point mass particles. This allows for accretion to be easily turned on in future simulations if it is desired. As the cores interact with the gas very little through collisions, the sink particles are only allowed to interact with the gas particles gravitationally. The gravitational interaction between sinks and gas is a “softened” interaction, meaning that at large radii, the potential is very close to the true value, but when the distance between sink and gas particle is on the order of the “softening length”, the interaction is significantly less strong. The softening function can take a number of shapes. In PHANTOM, we have used either a Plummer potential (for the very first simulations we carried out) or the true potential softened with the cubic spline (for the more accurate simulations). The gravitational potential at a particular point in space due to a sink particle is given by:

$$\Phi(r, h) = -GM\phi(r, h_{\text{soft}}), \quad (2.16)$$




Figure 2.1: Comparison between the true potential (green), the Plummer softened potential (red) and the potential softened with a cubic spline (blue). These were calculated with a mass of $0.6 M_{\odot}$ and a softening length of $3 R_{\odot}$.

where $\phi(r, h_{\text{soft}})$ is the softening kernel. The Plummer softening kernel is given by:

$$\phi(r, h_{\text{soft}}) = (r^2 + h_{\text{soft}}^2)^{-1/2},$$

allowing us to write the Plummer potential as:

$$\Phi(r, h_{\text{soft}}) = -\frac{GM}{\sqrt{r^2 + h_{\text{soft}}^2}}, \quad (2.17)$$

where M and h_{soft} are the mass and softening length of the sink particle, respectively. The softening kernel based on the cubic spline is found by the relation:

$$W(r, h_{\text{soft}}) = \frac{1}{4\pi r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial \phi(r, h_{\text{soft}})}{\partial r} \right), \quad (2.18)$$

where $W(r, h_{\text{soft}})$ is the cubic spline given in equation 2.3. Solving the equation for the softening kernel gives

$$\phi(r, h_{\text{soft}}) = \frac{1}{h_{\text{soft}}} \begin{cases} \frac{q^5}{10} - \frac{3q^4}{10} + \frac{2q^2}{3} - \frac{7}{5} & \text{for } 0 \leq q \leq 1 \\ -\frac{q^5}{30} + \frac{3q^4}{10} - q^3 + \frac{4q^2}{3} - \frac{8}{5} + \frac{1}{15q} & \text{for } 1 \leq q \leq 2 \\ -\frac{1}{q} & \text{for } q > 2 \end{cases} \quad (2.19)$$

where $q \equiv \frac{r}{h_{\text{soft}}}$ as before. Hence, the potential due to a sink can be found with the use of equation 2.16. The Plummer potential was introduced initially for the preliminary simulations as it is faster to implement and run. However, it only asymptotically approaches the true potential at large r , while the cubic spline softened potential becomes the theoretical potential outside of $r = 2h_{\text{soft}}$. These two potential shapes can be seen in Fig. 2.1.

A typical feature of SPH codes is to allow the timestep of a simulations to vary based upon some constraints. This ensures that the accelerations and velocities of particles are calculated in sufficiently small increments, while also decreasing the computational cost. There are several criteria which can dictate the timestep of a particle in PHANTOM. The primary restriction is the commonly used Courant condition:

$$\Delta t_C = C_{\text{cour}} \frac{h}{v_{\text{sig}}} \quad (2.20)$$

where $C_{\text{cour}} = 0.3$ is the default Courant factor and v_{sig} is the maximum signal speed over all neighbouring particles. Enforcing this condition ensures that compression waves and similar features will take at least three timesteps to move through a particle, therefore the features are sufficiently resolved in time. A couple more timestep conditions that are used in our simulations are based upon forces on the particles:

$$\Delta t_f^a = C_{\text{force}} \sqrt{\frac{h_a}{|\mathbf{a}_a|}} \quad (2.21)$$

$$\Delta t_{\text{sink-gas}}^a = C_{\text{force}} \sqrt{\frac{h_a}{|\mathbf{a}_{\text{sink-gas}}|}} \quad (2.22)$$

where $C_{\text{force}} = 0.25$ is the default value, h_a is the smoothing length of particle a and $|\mathbf{a}_a|$ is the magnitude of the acceleration on particle a . These are known as the ‘force condition’ and the sink-gas force condition respectively. There are some other timestep conditions, but they are not particularly important for our simulations. The true timestep for a particular particle is then the minimum of these conditions:

$$\Delta t_a = \min(\Delta t_C, \Delta t_f^a, \Delta t_{\text{sink-gas}}^a) \quad (2.23)$$

The minimum value of Δt_a over all particles a in the simulation will generally be set as the timestep of the simulation. If an individual timestepping scheme is adopted, however, the timestep that a particle uses will be based upon their Δt_a . The idea behind independent timesteps is very simple. A set of timestep ‘bins’ is created by the code, with the largest having an arbitrary maximum time, Δt_{max} , as an input to the code, while the other bins are created in factor-of-two decrements from the bin above them. The maximum number of these bins is decided at the beginning of the simulation. Particles are then sorted into the timestep bin that is closest to, and smaller than, their calculated Δt_a . The factor of two difference between successive timestep bins forces the particles to sync up at least every Δt_{max} . This sort of scheme has costs in the conservation of energy and angular momentum, but if it is properly implemented, it can speed up simulations by an order of magnitude (Price et al., 2015).

Simulations in this paper were run on 16 or 32 core machines, as PHANTOM currently only has OpenMP implemented, rather than MPI. Approximate wall clock times (the exact times were unfortunately not recorded) and number of cores for each simulations can be found in Table A.1 (Appendix A).

3

Benchmark Simulations

In the previous chapter, we outlined the SPH method, gave a brief description of the advantages and disadvantages of SPH and grid codes, and introduced the SPH code that we will be using in our simulations, PHANTOM. Up until this point we have been setting the scene, but now we will begin to discuss the simulations of Passy et al. (2012) in greater detail. Furthermore, we will cover in detail the benchmarking of our own simulations, as well as some of the issues that we have encountered so far.

The star stabilisation process in PHANTOM, described in Section 3.1, was initially performed by Roberto Iaconi. This work helped define some of the initial stabilisation parameters that served as a platform for our simulations.

3.1 Repeating the simulations of Passy et al. (2012) with PHANTOM

Before being able to perform any kind of rigorous science, it was necessary to have some kind of benchmark against which to set any future results. Passy et al. (2012) carried out common envelope simulations using both three dimensional grid and SPH codes. The comparison between the two methods is useful to ensure that key outputs are the same between the two types of code, giving a greater confidence that we understand the physics at play; if features are common to the outputs of both types of code, then they are more likely to be physical in nature, rather than an artifact dependent on the numerical method. The results of the simulations carried out by Passy et al. (2012) have therefore been used to calibrate PHANTOM.

Passy et al. (2012) ran SPH and grid simulations, using the codes *SNSPH* and *Enzo*, respectively. The SPH simulation had 5×10^5 particles, and the *Enzo* simulations were carried out with two uniform grid resolutions of 128^3 and 256^3 cells. A one-dimensional model star was evolved to the red giant branch (RGB) with the stellar evolution code EVOL (Herwig, 2000), with the density and pressure profiles mapped into three dimensions in the two

different codes. The *SNSPH* star had the initial setup modified so as to be in hydrostatic equilibrium, while the *Enzo* stars were allowed to relax in the computational domain. This difference in the setup caused a slight disparity in the radii of the stars, where the *SNSPH* star had a radius of $83 R_{\odot}$, while the low and high resolution *Enzo* stars had radii of $91 R_{\odot}$ and $85 R_{\odot}$, respectively. The cores of these giant stars were approximated to be point mass particles, in both codes, primarily because the resolution required to simulate the core properly is too computationally expensive. Having so much mass packed into a small volume will also generally decrease the timestep of the simulation. The companion star was also modelled as a point mass, as it is a main sequence star of less than $1 M_{\odot}$ and so it would have a radius of less than $1 R_{\odot}$, which is very small compared to the primary and similar to, or smaller than, the resolution of the simulations. The simulations of Passy et al. (2012) had the companion star placed at the surface of the primary at the beginning of the simulation. This initial setup is useful as a way to approximate the plunge-in stage of the common envelope. While they simulated systems with a range of companion masses, we will be focusing only on the case with a secondary star mass of $0.6 M_{\odot}$.

It is worth noting as well that, in their SPH simulations, Passy et al. (2012) used point mass particles, rather than sink particles. These point mass particles had both a “gravitational mass”, which was used to calculate the gravitational potential of the particle, and a vastly different “SPH mass”, which allowed the point mass to exert a pressure on the surrounding gas. The SPH mass of their primary core was chosen so as to have the star in hydrostatic equilibrium, while the companion’s SPH mass was set to zero, so that it would act as an n -body particle. The core of the primary had a static smoothing length of $0.1 R_{\odot}$, but it is not clear if any particular value was assigned to the companion. This value of $0.1 R_{\odot}$ can be thought of as being analogous to the softening length used in our two sink particles, discussed in Chapter 2.

We used the same one dimensional stellar profile that was used by Passy et al. (2012) for our primary star. The stellar evolution codes used to produce such profiles use the equations of stellar structure along with an understanding of radiative processes, nuclear physics and the equation of state to follow the evolution of a star of arbitrary initial mass. The 1D star was evolved from a $0.88 M_{\odot}$ main sequence star to being on the RGB, at which point the radius reached $83 R_{\odot}$. The 1D data we used consisted of density, pressure and temperature as a function of radius, which were mapped into three dimensions in the PHANTOM SPH code. For our simulations, we used the adiabatic equation of state

$$P = (\gamma - 1)\rho u \quad (3.1)$$

where P is pressure, $\gamma = \frac{5}{3}$ is the adiabatic index, ρ is density and u is specific internal energy.

This ‘star’ was relaxed in the three dimensional computational domain by leaving it unperturbed and damping any velocities for approximately 80 days (or ~ 4 dynamical times), such that there were no longer oscillations in the star. At this point, the star was able to be evolved with undamped velocities for at least another 4 dynamical times, within which time it remained essentially static. In the first timestep of the stabilisation, material within a radius of $0.03 R_{\odot}$ from the point of maximum density (i.e., the centre) was artificially “sunk” into a particle. This radius was chosen such that the “sink particle” would grow to a mass of $0.392 M_{\odot}$, which is the same as used by Passy et al. (2012). In the initial stellar profile, the star has an encapsulated mass of $0.392 M_{\odot}$ at approximately $0.11 R_{\odot}$ from the centre,

although $0.390 M_{\odot}$ is contained within the previously mentioned $0.03 R_{\odot}$. The discrepancy here allows for gas to be accreted onto this sink particle while the star is being stabilised. This radius is clearly quite small when compared to the softening length of $3 R_{\odot}$, however, a larger softening length decreases the time taken to run the simulations. A sink particle was also used in the simulation to represent a compact companion with a mass of $0.6 M_{\odot}$. The companion sink particle was not sunk from any existing SPH particles, instead, it was simply placed by the code. Both of the sink particles had softening lengths of $3 R_{\odot}$ (see Equation 2.19 and Fig. 2.1). Further, the sink particles were not allowed to accrete material during the simulation, though the use of sink particles instead of point mass particles allows accretion to be modelled when necessary. When setting up the first simulations, the sink particles employed a Plummer softened potential, however, this was soon replaced by softening the potential with a cubic spline kernel in an attempt to improve accuracy, though the results did not differ significantly. The difference between the softening length that we used and the core smoothing length used by Passy et al. (2012) is that our simulations were allowed to progress somewhat faster, and there appears to be only a minimal impact on the results. The companion star was placed at $100 R_{\odot}$ rather than on the surface of the primary. Initially, approximately 80 000 particles were used in the simulations. However, as we will show later by carrying out a convergence study, these simulations were not particularly accurate due to dramatically under-resolving the outside of the primary star as well as any ensuing regions of low density.

Finally, for the most part, our simulations used only a global timestep for the particles. As was discussed in Section 2.3, individual timesteps allow particles to be sorted into timestep bins, and have their accelerations updated when their timestep has passed. Although this is considerably more efficient, it can cause non-conservation when particles operating on small timesteps come into contact with those on longer timesteps, as is likely to occur in common envelope interactions. Our preliminary simulations employed individual timesteps, however, they displayed non-conservation of energy on the order of 10% over a period of ~ 1000 days. For this reason, a global timestep was used in all subsequent simulations, with energy non-conservation dramatically reduced to be on the level of 0.1%.

One of the main problems with most prior simulations of the common envelope interaction is that the envelope is lifted but remains primarily bound to the cores, a result that is at odds with observations (see Section 1.3). The amount of bound mass for our simulations was determined by checking that the sum of kinetic and gravitational potential energies for an SPH particle was positive, as follows:

$$\frac{1}{2}m_i v_i^2 + m_i \Phi_i > 0 \quad (3.2)$$

where m_i and v_i are the mass and velocity of particle i , respectively, while the gravitational potential on the particle, denoted by Φ_i , includes contributions from all other particles in the simulation as well as the sink particles (core and companion). There are alternative definitions of whether or not material is bound to the system. Ivanova et al. (2015) suggest that the addition of either the internal energy, u_i , or the enthalpy, $u_i + \frac{P_i}{\rho_i}$, should be used to determine whether mass is unbound. It is likely that exclusion of the internal energy will overestimate the amount of bound mass, however, the inclusion of internal energy or enthalpy will possibly underestimate the amount. For simplicity, then, we will stick with the definition in Equation 3.2.




Figure 3.1: Comparison of the bound envelope mass for simulations with different resolutions; blue is for the 8×10^4 particle simulation. The other three colours will be discussed in Section 3.2; green is for the 2.3×10^5 particle simulation, red is for the 1.1×10^6 particle simulation and cyan is for the 2.3×10^6 simulation.

We found that the bound mass follows a similar trend to that of Passy et al. (2012), wherein it drops quickly in the initial stages of the simulation, from a mass of $\sim 0.49 M_\odot$ to $0.44 M_\odot$, after which it levels off at ~ 100 days (see Fig. 3.1, blue line). However, the bound mass in our initial simulation began to quickly drop again at approximately $t = 900$ days (Fig. 3.1, blue line). As will be discussed further in Section 3.2, the time at which a second unbinding phase starts is dependent on the resolution of the simulation, with higher resolution simulations being able to run for longer without experiencing this secondary unbinding phase. Due to this, we suspect that this secondary unbinding is an artificial effect and the results of our simulations cannot be trusted after this unbinding occurs. This feature was not observed by Passy et al. (2012), possibly because their higher-resolution SPH simulations were not run for long enough to observe this resolution-dependent phenomenon (see Section 3.2).

3.1.1 Comparison of the evolution of the orbital separation

For the purposes of the comparison with Passy et al. (2012), we will henceforth discuss the values relating to a PHANTOM simulation with 2.3×10^5 particles, as this is the closest resolution we have to the 500 000 particle simulations undertaken by Passy et al. (2012). Further, it is able to be run for longer than 1000 days before displaying any artificial unbinding (see Fig. 3.1). Other resolutions will be discussed in Section 3.2, but for the most part, the evolution of orbital separation and energies between resolutions are fairly comparable, with the highest and lowest resolutions being within 6% of each other at 1000 days, showing the results to be reasonably well converged.

Possibly the easiest comparison that we can make between our simulation and that of Passy et al. (2012) is the orbital separation of the primary core and the companion. There will inevitably be a slight difference in results, as Passy et al. (2012) placed their companion




Figure 3.2: Comparison between the orbital separation of Passy et al. (2012) (top panel), and our simulations (bottom panel). The green line in the top panel corresponds to a simulation with a $0.6 M_{\odot}$ companion, the others have different masses.

on the surface of the primary, at $83 R_{\odot}$, while ours were placed initially at $100 R_{\odot}$. In spite of this, the evolution of the separations display very similar forms (Fig. 3.2), and by the end of 1000 days, they have reached an almost identical plateau value of $16 R_{\odot}$. This constitutes a reduction in the orbital separation of 81% in the Passy et al. (2012) simulation, and an 84% reduction in our simulations.

3.1.2 The inspiral timescale

What defines the beginning and the end of the dynamic inspiral phase is not particularly clear. While it is fairly simple to look at a plot of the orbital separation of two sink particles undergoing a common envelope interaction and immediately point out where dynamic inspiralling occurs, giving a more quantitative and physically meaningful definition of the moment at which the fast inspiral actually begins and ends is somewhat more tricky. A

quantative definition may assist us in comparing different common envelope simulations.

The timescale of the interaction is typically defined as $-\frac{a}{\dot{a}}$, where a is the orbital separation and \dot{a} is the time derivative of the separation (Sandquist et al., 1998). To calculate this quantity at a point in time, we averaged the separation over a 40 day period around the point in question. The timescale is essentially a metric for determining how quickly something is moving inward, and by averaging separations, we are trying to smooth out the short period oscillations. We present the timescale for our PHANTOM simulations in Fig. 3.3. Due to the fact that the simulation was initiated with the companion quite close to the primary surface, the timescale quickly drops to a minimum. It is this minimum that we use to determine where the fast inspiral occurs. We define the beginning and end of the fast inspiral to be at the times when the timescale is a factor of 5 greater than the minimum point. While arbitrary, this definition appears to do a fairly reasonable job of picking out both the beginning and the end of the dynamic inspiral phase, though further refinement may be necessary. By applying this method to our orbital evolution curves, we can later compare values at a set point after the dynamic inspiral occurs, regardless of the initial conditions, as can be seen from the vertical lines in Fig. 3.2 and Fig. 3.3.

3.1.3 Comparison of the energies and angular momenta

In Figs. 3.4 and 3.5, we show plots of the various energy components as a function of time. The three kinetic energy components, K_c for the kinetic energy of the cores, K_b for the KE of the bound material and K_u for the KE of the unbound gas, are calculated by summations of the form:

$$K_a = \sum_a \frac{1}{2} m_a |\mathbf{v}_a|^2 \quad (3.3)$$

where a can a sink particle, a bound gas particle or an unbound gas particle, m_a is its mass and \mathbf{v}_a its velocity. In the code, the calculation of the gravitational potential energy is also in




Figure 3.3: Timescale (defined as $-\frac{a}{\dot{a}}$) of the orbit using a 2.3×10^5 particle PHANTOM simulation. The rapid oscillations occurring later in the simulation are caused by a slight eccentricity of the orbit.




Figure 3.4: Plot of energies as calculated from the results of Passy et al. (2012) (figure 9 in that paper).




Figure 3.5: Comparison plot of energies from our simulation with 2.3×10^5 particles. In this plot: K_c is the total kinetic energy of the two cores, U_{tot} is the total internal energy of the gas, K_b is the kinetic energy of bound gas, K_u is the unbound kinetic energy, E_{env} is the envelope energy, ϕ_{env} is the potential energy of the envelope, E_c is the orbital energy of the cores, E_{tot} is the total energy of the system, ϕ_c is the potential energy between the sinks and ϕ_{tot} is the total potential energy.

three components, for gravitational interactions between the sink particles, $\phi_{\text{sink-sink}} (\equiv \phi_c)$, between sink particles and gas particles, $\phi_{\text{sink-gas}}$, and between pairs of gas particles, $\phi_{\text{gas-gas}}$. Similarly to the kinetic energies, we can write these components in the general form:

$$\phi_{a-b} = -G \sum_a \sum_b \frac{m_a m_b}{r_{ab}} \quad (3.4)$$

where G is the gravitational constant, a and b can both either be sink or gas particles for the calculation of the above components, and r_{ab} is the separation between the two particles. Therefore, we can write

$$\phi_{\text{tot}} = \phi_{\text{sink-sink}} + \phi_{\text{sink-gas}} + \phi_{\text{gas-gas}} \quad (3.5)$$

for the total potential energy of the simulation. Finally, the internal energy of the gas is stored by the particles in PHANTOM, so calculation of the total internal energy, U_{tot} , is performed with a straight summation over all particles in the simulation. The rest of the energies in Fig. 3.5 are able to be calculated as combinations of these components. The orbital energy of the sink particles is defined as the addition of the sink kinetic and potential energies:

$$E_c = K_c + \phi_c, \quad (3.6)$$

the envelope potential energy is defined as the potential energy of gas due to the gas particles and the primary sink particle:

$$\phi_{\text{env}} = \phi_{\text{gas-gas}} + \phi_{\text{sink 1-gas}} \quad (3.7)$$

where $\phi_{\text{sink 1-gas}}$ is the potential energy of the between the primary sink and the gas. Finally, the envelope energy is defined as the sum of the envelope potential energy, the total internal energy and the kinetic energy of the bound gas:

$$E_{\text{env}} = \phi_{\text{env}} + U_{\text{tot}} + K_b. \quad (3.8)$$

Fig. 3.4 shows the different energy components of the simulation carried out by Passy et al. (2012), while Fig. 3.5 details the energies obtained from our simulation. The curves in these plots are very similar, even down to the small oscillations. The most striking difference is the beat behaviour evident in the simulation of Passy et al. (2012), though lacking in our own energy curves. This is explained to be due to undersampling their simulation, with a mismatch between the oscillations of the energy curves and the sampling rate of the simulation (Passy et al., 2012). The core-core potential energy is less negative in our simulation because the two sink particles are placed slightly farther apart. Similarly, the envelope potential energy is initially less negative in our simulation because our star is slightly fluffier than that of Passy et al. (2012) (radius is approximately 11% greater in our simulations). The envelope potential energies differ by approximately 50% between the two simulations at 1000 days, likely due to the slightly different initial conditions as well as inherent code differences.

The kinetic energies are very similar between the two simulations. There is initially no unbound material in either simulation, so $K_u = 0$, however, by the end of our simulation the kinetic energy of the unbound mass is just starting to become greater than the bound kinetic energy, while this is not the case in the Passy et al. (2012) simulations. This effect is small, and is no longer present at higher resolutions. The differences in the kinetic energy of the bound mass are negligible. The variation in initial and final values, as well as the shapes of the curves, is again very likely to do with differing initial conditions and code differences.

Clearly, however, the largest change is in the internal energy of the gas, with our simulation having a value of U_{tot} at $t = 1000$ days that is 60% less than for Passy et al. (2012). However, while this percentage difference is large, if we compare the actual difference between initial and final values for the two simulations, the total internal energy changes by approximately 1.7×10^{46} ergs in both simulations, suggesting that the evolution of the two simulations is very similar, and that the initial conditions are very likely to be the reason for the differences.




Figure 3.6: Comparison of the evolution of angular momenta between Passy et al. (2012) (top panel, their Figure 8) and our 2.3×10^5 particle simulation (bottom panel). In this plot: J_{tot} is the total angular momentum, J_{orb} is the orbital angular momentum of the sink particles, J_{b} is the angular momentum carried by bound matter and J_{u} is for unbound gas. These curves include only the z-component of the angular momentum. Note that the limits on the y-axis are not the same.

The second set of values against which we can compare are the different contributions to the z-component of the angular momentum, J_z , which is by far the most dominant of the three orthogonal components. Further, if we are to be interested in material falling back to form a disc in the orbital plane, it makes sense that J_z is the most crucial component. The angular momentum was determined by taking the cross product of the particle radius with its momentum:

$$\mathbf{J} = \mathbf{r} \times m\mathbf{v} \quad (3.9)$$

where $\mathbf{J} = \begin{bmatrix} J_x \\ J_y \\ J_z \end{bmatrix}$, \mathbf{r} is the position vector, \mathbf{v} is the velocity vector and m is the mass of the particle. The z-component of the angular momentum then is just the third component of

the **J**. Fig. 3.6 offers a visual comparison of the different contributions to the total angular momentum, where J_b and J_u are the sums of J_z for all bound and unbound gas, respectively, and J_{orb} is the orbital angular momentum of the sink particles. Like the energy curves, the angular momentum curves are quite similar in shape, with the slight differences in the initial values of J_b and J_{orb} being consequences of our star being slightly fluffier and having the companion start at a larger orbital radius. It is also clear that the conservation of angular momentum is considerably better in our simulation.

3.2 Convergence tests

Passy et al. (2012) did not carry out a convergence test for their SPH simulations, although their grid simulations were carried out with two different resolutions. They carried out a theoretical SPH-grid resolution comparison and found that a 500 000 particle SPH simulation was slightly better resolved than a 256^3 unigrid simulation, and the results were very comparable.

To test for convergence of our PHANTOM common envelope simulations, we carried out 4 identical simulations with increasing resolutions (8×10^4 , 2.3×10^5 , 1.1×10^6 and 2.3×10^6 particles). Fig. 3.7 shows the highest resolution simulation at several snapshots in time to show details of the evolution. For each of the different resolutions, a new star had to be stabilised in the computational domain, which resulted in slightly different initial stellar parameters. The most prominent difference was that the primary star reached equilibrium at a slightly larger radius for the higher resolutions. For the 8×10^4 case, the star had an initial radius of $90 R_\odot$, whereas in the highest resolution, 2.3×10^6 particle simulation, the primary radius was $94 R_\odot$. The radii of these stars were determined by finding the volume equivalent radius (adapted from Nandez et al. (2014)) defined by:

$$R_V = \left(\frac{3}{4\pi} \sum_i \frac{m_i}{\rho_i} \right)^{\frac{1}{3}}, \quad (3.10)$$

where m_i and ρ_i are the mass and density, respectively, of particle i belonging to the star. These differences are likely caused by the pressure and gravity terms being more poorly resolved in the lower resolutions, hence leading to different equilibrium configurations. Further, it is possible that a star with greater resolution will simply resolve the low density atmosphere of the star better, and hence could result in a slightly increased radius. We do not expect these small differences in the initial radius to be the major cause of variations in the simulation outcomes.

We first compare the orbital separation evolution (Fig. 3.8). The shapes of the curves are very similar overall, but clearly only the 1.1×10^6 and 2.3×10^6 particle simulations are converged, as there are almost no differences between those two curves. At ~ 800 days, they both have a separation of $\sim 16.5 R_\odot$, with a variation of only 0.2%, and are only slightly out of phase. On the other hand, the 8×10^4 and 2.3×10^5 particle simulations have separations of 17.5 and $17.0 R_\odot$ respectively (a variation of 3%).

Secondly, we compare the amount of bound mass at the 500 days. Fig. 3.1 shows that all of the simulations initially unbind a similar amount of mass, and then level out for a period of time. It is the value of this plateau that is interesting when determining convergence.




Figure 3.7: Evolution of the 2.3×10^6 particle simulation at 0, 50, 150 and 1000 days. The left column shows a slice of the simulation in the x-y plane, and hence renders density. The right column shows the all the particles along the line of sight, therefore rendering the density integrated along the z-axis, the column density. The sink particles are plotted as cyan dots. Image created with SPLASH (Price, 2007).

Similarly to what was deduced from the separation, the 1.1×10^6 and 2.3×10^6 particle simulations appear to be almost identical at this time, having both unbound approximately $0.06 M_{\odot}$. Again, the simulations conducted with the lower two resolutions were clearly not well converged, having unbound 0.053 and $0.056 M_{\odot}$ for the 8×10^4 and 2.3×10^5 particle simulations, respectively.

The bound mass calculations, however, yielded a more curious result; they showed, as




Figure 3.8: Comparison of the orbital separation between the two sink particles; blue is for the 8×10^4 particle simulation, red is for the 2.3×10^5 particle simulation, magenta is for the 1.1×10^6 particle simulation and green is for the 2.3×10^6 particle simulation. Inset: details of the different resolutions at ~ 800 days.

noted in Section 3.1, a resolution-dependent rapid unbinding phase was observed in all simulations. Fig. 3.1 shows that for higher resolutions, this secondary unbinding phase occurs at later times. Further, the gradient of the decline in bound mass is significantly greater in the lower resolution cases. In general, resolution dependent results indicate artificial features, so we cannot evolve these simulations for longer than the time to begin unbinding. Fig. 3.1 also shows that for a resolution of 5×10^5 particles, as used by Passy et al. (2012), the secondary unbinding occurs much later than 1000 days, and so it is not surprising that it was not picked up by their simulations.

This artificial unbinding has been the cause of much discussion and various attempts to solve the problem. Although we have been, thus far, unsuccessful, we think we have correctly identified the problem and we are carrying out tests to determine a solution. Only when this has been accomplished can we continue the simulations to study the effect of the fallback phase.

Firstly, we can see clearly that, as the simulation progresses, the volume immediately around the orbit of the sinks is progressively evacuated, leading to a local decrease in density. This effect can be seen in Fig. 3.9, which shows the density of particles for three different resolutions at 1000 days. As can clearly be seen, there is an evacuation of particles from a volume around the two sink particles. At some point, which is earlier in lower resolution simulations, a density limit is reached below which a catastrophic, runaway unbinding takes place. As the density decreases in the area around the sink particles, an increase in smoothing length of the SPH particles also occurs. This can lead to SPH particles falling in too close to a sink particle, and consequently being flung on a ballistic trajectory out of the system.

To prevent this “flinging”, we modified the code to allow the softening length of the sink particles to change:

$$h_{\text{soft}} = \max(h_{\text{soft,c}}, h), \quad (3.11)$$




Figure 3.9: Comparison density of particles around the sinks for simulations with different resolutions; the left-hand panel is the 8×10^4 particle simulation, the middle panel is the 2.3×10^5 particle simulation, and the right-hand panel is the 1.1×10^6 particle simulation. Image created with SPLASH (Price, 2007).




Figure 3.10: Comparison of the bound mass for the simulations with constant $h_{\text{soft}} = 3 R_{\odot}$ (green), $\max(3, h) R_{\odot}$ (blue) and $\max(6, h) R_{\odot}$ (red).

where $h_{\text{soft},c} = 3 R_{\odot}$ is the original softening length of the sink particles, and h is the variable smoothing length of the particle with which the sink is interacting. This typically would lessen the gravitational interaction between low density particles in the region close to sink particles, while all other sink-gas gravitational interactions in the simulation would remain the same. The result of this addition to the code was a slight increase in the time at which the unbinding took place, but otherwise, the shape of the unbinding curves remained almost identical (see Fig. 3.10).

It seems plausible that by lessening the interaction between sink and gas particles, there would be fewer gas particles to gain a ballistic trajectory. Hence, the next step was to try increasing the value of $h_{\text{soft},c}$ to be $6 R_{\odot}$ instead. Therefore, the gas particles would typically

be less bound to the sink particles, but would also have a lesser interaction. This simulation was run with limited success. While Fig. 3.10 shows that the bound mass was greater at all points in the simulation than for $h_{\text{soft,c}} = 3 R_{\odot}$, the artificial unbinding was still present and started earlier.

4

The Pre-Common Envelope Phase

In the previous chapter, we benchmarked the SPH code, PHANTOM, against the results of Passy et al. (2012) with a good degree of agreement between our simulations and theirs. In this chapter we will investigate how the simulated interaction is affected when we increase the initial orbital separation of the sink particles, looking primarily at a system that starts when on the verge of beginning Roche-lobe overflow. We will also briefly look at a system with an even larger initial orbit.

Further, we have carried out some relatively underdeveloped simulations in the corotating frame of reference. In such a frame it is simpler to stabilise the binary and carry out a study of the pre-common envelope phase, where we can determine how small instabilities influence the final outcome. In the future it will become possible to simulate two full stars with this method, as done preliminarily by (Nandez et al., 2014).

4.1 The Roche-lobe overflow phase

An area of concern in common envelope interaction simulations is the phase leading up to the dynamic inspiral. It is not known how different initial conditions, such as starting the simulations before Roche-lobe overflow (RLOF), may influence the final outcome of the interaction. This phase begins when the envelope of the primary begins to overflow its Roche lobe, leading to a period of mass transfer. This typically occurs when the radius of the primary increases as it becomes a red giant, eventually becoming larger than the Roche lobe radius, or if the companion is brought closer to the primary by tidal forces. There is no clear agreement on when this mass transfer phase is stable, and thus long-lived, or unstable, quickly reducing the orbital separation and leading to the common envelope inspiral (Ivanova et al., 2013). In this chapter we describe the outcomes of simulations that have been performed to study this stage of Roche-lobe overflow, particularly with respect to how it affects the final parameters of the binaries.

For the sake of speed, the primary simulation discussed in this chapter was carried out




Figure 4.1: Density slices in the orbital plane of the simulation starting at $218 R_{\odot}$, showing the evolution through time of the system. Dynamic inspiral occurs at approximately 8 years. The sink particles are plotted as cyan dots. Image created with SPLASH (Price, 2007).

with 2.3×10^5 particles and compared with the similar resolution simulation from Chapter 3. The lower resolution did not appear to cause particularly large changes (see Section 3.2). We realise that the differing initial conditions of the current simulation could dictate a different convergence behaviour, but a rigorous convergence test will be left for future work.

Setting up the primary star for this simulation was identical to the method described in

Section 3.1. In order to begin the simulation just before RLOF, it was necessary to place the companion such that the Roche lobe radius of the primary was just larger than its stellar radius. The Roche lobe radius can be approximated, for circular orbits, with the equation (Eggleton, 1983):

$$r_{RL} = \frac{0.49q^{2/3}}{0.6q^{2/3} + \ln(1 + q^{1/3})} a, \quad (4.1)$$

where r_{RL} is the Roche lobe radius of the donor star, a is the orbital separation and q is the ratio of the masses of the donor star and its companion ($q \equiv \frac{M_1}{M_2}$ in our notation). Hence, with a primary star of radius $r_1 = 90 R_\odot$ and mass $M_1 = 0.88 M_\odot$, and a companion of mass $M_2 = 0.6 M_\odot$, we can determine that RLOF will begin approximately when the separation, a , is such that $r_1 \approx r_{RL} \approx 218 R_\odot$.

The companion star was therefore placed at this distance to simulate a binary system on the cusp of initiating Roche-lobe overflow. The stars were given velocities such that they would be in a circular orbit, and the simulation was evolved until the sink particles reached a stable, compact orbit. Fig. 4.1 shows snapshots of the simulation, where RLOF is clearly visible in four of the panels. In Fig. 4.2 we show the orbital evolution, where the timescale (Fig. 4.3) has been used to divide the common envelope interaction into distinct phases. Fig. 4.2 places the beginning of the fast inspiral at a separation of $110 R_\odot$, 8.4 years after RLOF begins. Using our definition of the dynamical inspiral (see Section 3.1.2), this phase lasts for about half a year, within which time the orbital separation decreases down to $31 R_\odot$.

In the following, I will refer to the 2.3×10^5 particle simulation described in Chapter 3, starting at $100 R_\odot$, as *100Sep*, and the simulation described here, starting at $218 R_\odot$, as *218Sep*. First, we compare how the separations evolve in time. In Fig. 4.2 we compare the orbital evolution of the two simulations, by shifting the *100Sep* simulation by 3096 days such that its beginning coincides with the time at which the *218Sep* simulation has reached an orbital separation of $100 R_\odot$. The orbital evolution of the fast inspiral phase is quite similar between the two simulations. As might be expected, due to its larger initial angular momentum budget, the *218Sep* simulation ends up in a somewhat wider orbit of $20 R_\odot$, compared to the final separation of $16 R_\odot$ for the *100Sep* simulation. Therefore the separation of the two final orbits differ by approximately 30%. We can see that both simulations develop some eccentricity, with the *100Sep* simulation having an eccentricity of $e = 0.22$ at $t = 1000$ days (4096 days in Fig. 4.2), while the *218Sep* simulation has a smaller eccentricity of $e = 0.10$ at $t = 4096$ days. This eccentricity is likely due to the three-dimensional distribution of gas. However, the fact that the *100Sep* simulation has a larger eccentricity is probably due to the fact that the companion was initially on the surface of the primary resulting in a larger initial perturbation.

We can also appreciate the difference between the simulations by looking at the distribution of the gas and the associated velocity field (see Fig. 4.4). At $a = 100 R_\odot$, the *100Sep* simulation is unperturbed. However, by the time it reaches $a = 100 R_\odot$, the *218Sep* simulation presents a very different gas distribution. At that time, the gas particles in the *218Sep* simulation form a diffuse, extended envelope around the central binary due to interactions in the RLOF phase. It is to be expected that having less realistic starting parameters, such as those in the *100Sep* simulation, should decrease the validity of the results.

Fig. 4.5 shows a comparison between the evolution of the various energy components




Figure 4.2: Orbital separation of the two cores for the *218Sep* simulation, beginning from Roche lobe overflow (blue), and the *100Sep* simulation (green), which was shifted by 3096 days. The beginning ($t = 3067$ days) and end ($t = 3235$ days) of the dynamical inspiral phase (vertical lines) was again determined by using the timescale (shown in Fig. 4.3).




Figure 4.3: Timescale (defined as $-\frac{a}{\dot{a}}$) of the *218Sep* simulation. Note that a truncated x -axis is used here for clarity.

in the *100Sep* and *218Sep* simulations. The different energy components in Fig. 4.5 are described in detail in Section 3.1.3. The largest differences between the simulations in Fig. 4.5 lie mostly with the initial values of the internal energy and the potential energy of the gas (which show up in ϕ_{env} and ϕ_{tot} curves), and final values of the kinetic and potential energies of the cores. We remind the reader that the *218Sep* simulation “initial” values are taken at $t = 3096$ days when the sink particles have reached a separation of $100 R_{\odot}$. At this time, the star has been sufficiently distorted so as to be different from the *100Sep* simulation star at $t = 0$. The final value of K_c is lower for the *218Sep* simulation than for the *100Sep* simulation, as the cores that are more widely spaced will have a lower orbital velocity. The difference in ϕ_c can be explained in a similar manner, as the cores end up a $\sim 30\%$ larger




Figure 4.4: A density slice and velocity field (arrows) for the *218Sep* simulation when the sink particles have reached an orbital separation of $100 R_{\odot}$ ($t = 3096$ days), the initial separation for the *100Sep* simulation. The boundary here is $800 R_{\odot}$ on each side. Image created with SPLASH (Price, 2007).

orbit for the *218Sep* simulation, leading to a significantly less negative ϕ_c . On the other hand, the evolution of the gas particle quantities, U_{tot} , K_b , K_u , E_{env} and ϕ_{env} , are actually quite similar between the two simulations, particularly in their final values. This is because the difference in energies between the two simulations is carried by the sink particles, and the gas tends to plateau at around the same values.

The angular momentum of the *218Sep* simulation is particularly likely to vary significantly from the *100Sep* simulation, as its angular momentum budget is considerably greater due to the larger initial separation. The value of J_{tot} for the *218Sep* simulation, as can be seen in Fig. 4.6, is $\sim 40\%$ greater than for the *100Sep* simulation. The angular momentum of the bound mass reaches a peak at ~ 3200 days, a feature that is not seen in the *100Sep* simulation. After the peak it begins to decrease, while J_u begins to increase. This peak in the bound angular momentum roughly coincides with the time at which the fast inspiral occurs, at which point the envelope quickly begins to become unbound (Fig. 4.7). It is therefore likely that J_b drops for the simple reason that there is less bound mass. It is worth noting that each of the angular momentum components of the *218Sep* simulation is always greater than the respective components of the *100Sep* simulation. This is not surprising, given the initial conditions of the simulation, but it points to the envelope in the *218Sep* simulation likely being less bound at the end of the simulation, with the cores still retaining plenty of orbital angular momentum that can work towards the unbinding of the envelope.

Fig. 4.7 shows that the mass of the bound envelope for the *218Sep* simulation decreases slowly though the period of RLOF until the fast inspiral begins at 3067 days. At this point, the sink particles rapidly decrease their separation, transferring energy and angular momentum to the envelope. A larger angular momentum budget causes slightly more mass to be unbound over the course of the interaction; the *100Sep* simulation unbound $\sim 12\%$ of the mass of the envelope, while the *218Sep* simulation unbound $\sim 15\%$ of the envelope mass.




Figure 4.5: Comparison of the energies for the simulations *100Sep* (top panel) and *218Sep* (bottom panel). Note that the *218Sep* simulation has been plotted for 1000 days from $t = 3096$ days, so as to compare the two sets of curves. In this plot: K_c is the total kinetic energy of the two cores, U_{tot} is the total internal energy of the gas, K_b is the kinetic energy of bound gas, K_u is the unbound kinetic energy, E_{env} is the envelope energy, defined as $E_{\text{env}} = \phi_{\text{env}} + U_{\text{tot}} + K_b$, where ϕ_{env} is the potential energy of the envelope; E_c is the orbital energy of the cores, E_{tot} is the total energy of the system, ϕ_c is the potential energy between the sinks and ϕ_{tot} is the total potential energy.

Note again that the envelope once again begins to become catastrophically unbound at approximately ~ 4800 days (see Section 3.2). It can be seen fairly clearly in Fig. 4.7 that this is relatively later than in the *100Sep* simulation, likely because the density near the core drops more slowly in the *218Sep* simulation.

4.1.1 The corotating frame

In nature, companions out to $\sim 2\text{--}3$ stellar radii will be tidally captured. Before the rapid infall starts, the system will be almost or completely synchronised, resulting in a spinning giant (Tassoul and Tassoul, 1992). Evolving a binary system in the corotating frame, wherein the primary's spin is tidally locked to the orbit, allows for a simple stabilisation of a corotating




Figure 4.6: Evolution of the different components of the angular momentum in the *100Sep* (top panel) and *218Sep* (bottom panel) simulations. Note that the *218Sep* simulation has been plotted for 1000 days from $t = 3096$ days, so as to compare the two sets of curves. In these plots: J_{tot} is the total angular momentum, J_{orb} is the orbital angular momentum of the sink particles, J_b is the angular momentum carried by bound matter and J_u is for unbound gas.

primary star. A rotating star cannot easily be stabilised in the inertial reference frame because our stabilisation is partly performed by damping the gas velocities at each timestep. However, by choosing a rotating frame of reference such that the primary star is stationary, damping spurious velocities within the rotating star is simplified. In SPH, it is relatively straightforward to implement this reference frame, and we have done so.

The corotating frame was implemented by first calculating the angular velocity of the system with the standard formula (adapted from Paczynski (1971))

$$\Omega = \sqrt{\frac{G(M_1 + M_2)}{a^3}}, \quad (4.2)$$

From here, Ω was taken to be the angular velocity of the corotating frame. Using this value,




Figure 4.7: Evolution of the bound envelope mass in the *218Sep* simulation (blue) and the *100Sep* simulation (red). The *100Sep* simulation has been shifted by 3096 days.

it is possible to transform velocities between the original and rotating frames using:

$$\mathbf{v}_r = \mathbf{v}_i - \boldsymbol{\Omega} \times \mathbf{r}, \quad (4.3)$$

where \mathbf{v}_r is the velocity in the rotating frame, \mathbf{v}_i is the velocity in the inertial frame, \mathbf{r} is the position vector and

$$\boldsymbol{\Omega} = \begin{bmatrix} 0 \\ 0 \\ \Omega \end{bmatrix}.$$

Taking the time derivative of the velocity gives the acceleration in the rotating reference frame:

$$\mathbf{a}_r = \mathbf{a}_i - 2\boldsymbol{\Omega} \times \mathbf{v}_r - \boldsymbol{\Omega} \times (\boldsymbol{\Omega} \times \mathbf{r}), \quad (4.4)$$

where the accelerations due to the ‘fictitious’ centrifugal and Coriolis forces are given by $\mathbf{a}_{\text{cen}} = -2\boldsymbol{\Omega} \times \mathbf{v}_r$ and $\mathbf{a}_{\text{cor}} = -\boldsymbol{\Omega} \times (\boldsymbol{\Omega} \times \mathbf{r})$, respectively.

Only very preliminary work has been undertaken in the corotating frame for the purpose of testing that the analysis routines were able to properly calculate all quantities in a non-inertial frame. The system was set up with the methods described in Section 3.1, again with a separation of $218 R_\odot$. We will call this simulation *218SepCor*. While, the system was not stabilised in the corotating frame, the long orbital decay in Fig. 4.8 suggests that the star was quite stable anyway, so a comparison with the simulation carried out in the inertial frame is not complete. We nonetheless give a brief review of how having a corotating primary star changes the outcomes of the simulation. We may expect some changes, because when the star is tidally locked with the orbit, the system has a higher total angular momentum than when the star has no spin.

The orbital separation is reduced when angular momentum is transferred from the orbit to the envelope of the primary, which can be achieved through a couple of methods; tidal interactions and gravitational drag. In our simulations, the *218SepCor* simulation would suffer




Figure 4.8: Evolution of the orbital separation in the *218SepCor* simulation (blue) and the *218Sep* simulation (red).

a weaker tide than for the *218Sep* simulation. For a system not in corotation, the tides (which will be discussed in slightly more depth in Section 4.2) raise bulges in the primary, which are slightly out of alignment with the companion. This works to drag the primary into corotation by increasing its spin angular momentum, requiring a reduction of the orbital angular momentum by decreasing the orbital separation. However, when a system is already in corotation, such as in the *218SepCor* simulation, these bulges will have very little effect on the orbital evolution, so the decline in separation will take much longer. Further, in a corotating star, the velocity contrast between the companion and the stellar layers is noticeably reduced, making the strength of the gravitational friction much lower. Thus, the orbital separation is reduced far more slowly in the *218SepCor* simulation than for the *218Sep* simulation. The orbital separation evolution curves for the *218Sep* and *218SepCor* simulations are compared in Fig. 4.8 where we can see that both simulations plateau at approximately the same final orbital separation of $20 R_{\odot}$.

The mass of the bound envelope drops to a lower value in the *218SepCor* simulation ($0.393 M_{\odot}$) than for the *218Sep* simulation ($0.416 M_{\odot}$), as the envelope of the primary star is initially much less bound in the *218SepCor* simulation (Fig. 4.9). This is a direct consequence of the gas having a larger kinetic energy than the *218Sep* simulation. It should be noted that, due to the fact that this simulation was not stabilised in the rotating frame, the effect here is likely to be somewhat less pronounced. Stabilisation should lead to the envelope of the primary increasing slightly in radius, thus having a less negative potential energy and increased kinetic energy if it is artificially kept in corotation.

During the relatively stable phase preceding the rapid infall, there is quite a large amount of unbinding. The *100Sep*, *218Sep* and *218SepCor* simulations are all progressively more stable, and result in progressively more unbound gas. Once this pre-infall phase has taken place, the envelope distribution will be quite altered and finding an explanation of why there is more or less unbinding becomes more complex.




Figure 4.9: Evolution of the bound envelope mass in the *218SepCor* simulation (blue) and the *218Sep* simulation (red). Note: the downturn in the red curve at ~ 5000 days is due to the artificial unbinding we have discussed in Section 3.2.

4.2 The pre-Roche lobe overflow phase

Before a binary system goes through Roche-lobe overflow and a common envelope interaction, it experiences a phase of tidal interaction. Tidal interactions have been studied in great depth by Zahn (1977, 1989). Of primary importance for binary systems in the pre-RLOF phase is the equilibrium tide, which is essentially characterised by a deformation of a star in hydrostatic equilibrium when exposed to an external gravitational potential. The result is that two opposite bulges form on the primary star aligned with the direction of the companion (the companion forms similar bulges along the direction to primary, but due to its compact nature, the effect is negligible). In the absence of dissipative forces, these bulges would lie directly along the line connecting the centres of mass of the two stars. However, convective dissipation is present in giant stars so the bulge may precede or, more typically for slow rotating giant stars, lag the line between the centres. This results in the slowing down of the stars, decreasing the orbital separation while spinning up the primary star (Hurley et al., 2002).

Simulations are generally unable to model tides correctly as the primary form of dissipation, arising from turbulent viscosity in stellar convection, is not properly described. On the other hand, artificial viscosity, which was discussed in Section 2.1, does result in dissipation that can qualitatively act as tides.

We have run a simulation of the evolution of a binary system in the tidal interaction phase, that is, before Roche-lobe overflow, using the same setup as the *218Sep* simulation with an initial separation of $300 R_{\odot}$. The orbital separation of this simulation changes very little over the course of ~ 20 years (see Fig. 4.10). A very small amount of mass transfer takes place in the system, starting at approximately $t = 1200$ days, but far too little to cause the orbital decay seen. We suspect that a combination of the numerical viscosity and tidal torque from oscillations in the primary envelope is the cause of this “tidal” decay. The decay




Figure 4.10: Evolution of the orbital separation for a simulation with $300 R_{\odot}$ initial separation.

timescale, while slow for our simulation, is still quite fast compared to the 10^4 to 10^5 year timescales expected in nature (Mustill and Villaver, 2012). Due to the long decay timescale of our simulations (Fig. 4.10), we have been unable to continue this simulation for the length of time needed to compare results with other simulations, and we reserve this task for future work.

In Fig. 4.10, the separation decreases over 9000 days by $\sim 3\%$. At the same time, the system develops an eccentricity of ~ 0.2 . It is unclear what this eccentricity is due to, as it is larger than the eccentricities that form in the early stages of previous simulations (which are on the order of 10^{-3} for both the *218Sep* and *218SepCor* simulations). The reason for this orbital decay and increase in eccentricity is the sudden introduction of the companion near the stable primary. This excites an oscillation in the primary which exerts a variable tidal force, the net effect of which is a shortening of the orbital separation and an increase in eccentricity. Although this is an artificial effect, it serves to reduce the orbital separation within reasonable computational times. Once mass transfer begins, the orbital decay will be driven by much stronger drag forces.

Fallback of Gas During the Common Envelope Interaction

The study of fallback material is one of the ultimate goals of this research and one of the reasons for continuing to use SPH for the common envelope problem. As discussed in Chapter 3, understanding the impact of gas fallback on the common envelope outcome has been somewhat stymied by the inability of our simulations to be trusted beyond an unforeseen, resolution-dependent unbinding phase (see Section 3.2 and Fig. 3.1). However, we can search for fallback of material before the unphysical unbinding phase, as well as compare the fallback seen in simulations with differing initial conditions.

In addition, the recent work of Nandez et al. (2015) has demonstrated the potential importance of including recombination energy in simulations. If they are right and substantially more gas is unbound, then that source of energy must be included in all simulations. On the other hand, even with the recombination energy, some gas fallback is likely. Ahead of carrying out our own testing with recombination energy, we here simply consider the behaviour of the gas ejected during the dynamic inspiral phase, which will have a velocity range straddling the local escape velocity.

5.1 Examination of fallback gas in the simulation with a $100 R_{\odot}$ initial separation

The simplest method with which we can determine if material is falling back towards the central binary system is examining the radial component of the velocity of SPH particles at different times in the simulations. This quantity can be written as:

$$v_r = \frac{\mathbf{v} \cdot \mathbf{r}}{|\mathbf{r}|}, \quad (5.1)$$

in Cartesian coordinates, with \mathbf{r} and \mathbf{v} being the vectors of position and velocity, respectively, for an SPH particle. This quantity defines the magnitude of the velocity in the direction of the

origin, which, in our simulations, has been set at the centre of mass. Naturally, then, fallback material will have a negative value of v_r . Note that a negative radial velocity component is not necessarily indicative of fallback gas in the vicinity of the sink particles, because the primary star (i.e. the gas envelope and its core) moves with respect to the centre of mass (i.e. the origin), therefore approximately half of the gas acquires a periodic negative v_r , particularly at the beginning of the simulation when the star is undisturbed. While this particular issue may be abated by defining the radial velocity relative to the primary's core, this would introduce issues later in the simulation, wherein it would likely appear that expanded layers of the envelope with an overall low velocity were undergoing oscillatory radial motion.

In parallel, we also check if the gas is bound to the system. We use the specific energy (Equation 3.2) as a metric for whether or not a particle is bound. Fig. 5.1 shows the distribution of particles with different values of specific energy (top panels) and radial velocity (middle and bottom panels). The left-hand panels of Fig. 5.1 show that, at a time of 100 days, there are relatively few particles with a negative radial velocity. In comparison, at 1000 days (right-hand panels), a greater portion of the gas, spread over a larger volume of the simulation, has some kind of inward radial motion. At these moments in time, a large fraction of particles with outward motion are still bound to the system and hence will return at some point in the future, constituting a continuing fallback event. It is also possible for some of the gas to be unbound, but with a trajectory that will cause it to interact again with the dense inner region of the simulation.

While our simulations cannot be run for longer than ~ 1000 days, and thus we are unable to see fallback of large amounts of gas onto the central binary that would likely occur later (see below), there are clearly significant features that display an inward motion of gas in the lower panels of Fig. 5.1. Note that, while some of the gas with negative radial velocity in Fig. 5.1 (right-hand column, lower two panels) is due to the orbital motion of the gas, as explained above, this is not the case for all inbound gas. For example, note the inbound arch features on the right-hand side of the bottom right panel in Fig. 5.1.

We have based the discussion above on simulations run with 2.3×10^5 particles, because it would have taken a prohibitively long time to run simulations at the highest resolution with larger initial separations. However, in Fig. 5.1 we display results for the highest resolution of 2.3×10^6 particles, both because features are sharper at higher resolutions and because the simulation could be run for longer before the artificial unbinding (see Fig. 3.1 and Section 5.1), which has thus far prevented a more conclusive study of fallback gas.

Typically, in these simulations, the companion causes waves of density to ripple out through the envelope of the primary. In the initial stages of the interaction, a fraction of gas is unbound due to the pressure exerted by these density waves. However, by ~ 1000 days, these waves have diminished somewhat, and no more mass is being unbound. Instead, as the waves propagate through the gas, slightly evacuated regions are left behind, into which some of the surrounding gas will fall. This results in the spiral pattern appearance of the infalling material. This said, the radial motion of gas pushed out at below escape velocity is complex in the high density regions of the envelope, and fully analysing the particle trajectories will have to wait for a longer simulation run.




Figure 5.1: Slice in the orbital plane for the *100Sep* simulation with 2.3×10^6 particles. Left: $t = 100$ days. Right: $t = 1000$ days. Top: specific energy. Middle: radial velocity. Bottom: a rendering of negative radial velocity. The sink particles are plotted as black dots.

Tocknell et al. (2014) estimated the time it would take for bound material to return ballistically to the central binary after the end of the simulations of Passy et al. (2012). Assuming that all the bound gas would fall back towards the centre, Tocknell et al. (2014) calculated the radius of the circumbinary disc that would form from the infalling gas, by writing down the equations for the conservation of (specific) angular momentum and energy:

$$v_{\theta,h}h = v_{\theta,\text{disc}}r_{\text{disc}}, \quad (5.2)$$

where h is the average distance of the bound mass from the central binary at the end of the simulation, r_{disc} is the final radius of the circumbinary disc and $v_{\theta,h}$ and $v_{\theta,\text{disc}}$ are the tangential velocities of the bound mass and of the circumbinary disc, respectively. And:

$$\frac{1}{2}v_{\theta,h}^2 - \frac{GM}{h} = \frac{1}{2}v_{\theta,\text{disc}}^2 - \frac{GM}{r_{\text{disc}}}. \quad (5.3)$$

Rewriting Equation 5.2 for $v_{\theta,\text{disc}}$, substituting into Equation 5.3 and rearranging gives the quadratic equation:

$$\left(v_{\theta,h}^2 - \frac{2GM}{h}\right)r_{\text{disc}}^2 + 2GM r_{\text{disc}} - v_{\theta,h}^2 h^2 = 0, \quad (5.4)$$

which can be solved to give:

$$r_{\text{disc}} = \frac{v_{\theta,h}^2 h^2}{2GM - v_{\theta,h}^2 h}. \quad (5.5)$$

We can determine that, at 1000 days, the bound mass is distributed out to about $h = 2500 R_{\odot}$. We used a value of $v_{\theta,h}h = 1.7 \times 10^{18} \text{ cm}^2 \text{ s}^{-1}$ from our *100Sep* simulation, found dividing the bound angular momentum (Fig 3.6) by the mass of the bound envelope at $t = 200$ days (Fig. 3.1). Hence, dividing $v_{\theta,h}h$ by h , we obtain $v_{\theta,h} = 9.9 \times 10^4 \text{ cm s}^{-1}$, the radius of the circumbinary disc is estimated to be on the order of $11 R_{\odot}$, which is smaller than the binary separation, demonstrating that at least some material will re-interact with the binary. Tocknell et al. (2014) used the equation:

$$t = \pi \sqrt{\frac{\left(\frac{h+r_{\text{disc}}}{2}\right)^3}{GM}}, \quad (5.6)$$

to calculate the time it would take the bound gas to fall back. This yields a fallback timescale of approximately 6 years for our bound material. Note that this is a very approximate method of determining fallback time. Unfortunately, this timescale is larger than the time period of ~ 1000 days for which we are able to simulate at present. Even for the highest resolution runs (2.3×10^6 particles), we can only simulate for a maximum of 5 years, which is smaller than the time we calculated for a fallback. While this 6 year fallback timescale is just an approximation, Tocknell et al. (2014) also mentioned that more accurate calculations have revealed longer fallback times than their method produces. Therefore, it is not surprising that a full fallback event is not observed within our simulations.

5.2 Examination of fallback gas in the simulation with a $218 R_{\odot}$ initial separation

Interestingly, when we look at the material falling back in a simulation beginning with a separation of $218 R_{\odot}$, it appears that after the dynamic inspiral there is a short period where




Figure 5.2: Slice in the orbital plane for the *218Sep* simulation. Left: $t = 9$ years. Right: $t = 10$ years. Top: specific energy. Middle: radial velocity. Bottom: a rendering of negative radial velocity. The sink particles are plotted as black dots.

much mass starts to return to the central binary. This occurs within about a year of the end of the dynamic inspiral, as compared to the *100Sep* simulation, for which gas starts to fallback approximately 2.5 years after the end of the inspiral phase. Although it may previously have made sense to compare these simulations at similar times after the dynamic inspiral, in this case we are looking only for the fallback of gas. For the *100Sep* simulations, the times of 100 and 1000 days were chosen to emphasize the early and maximal distributions of fallback gas. Recalling that the orbital separation of the *218Sep* is $100 R_{\odot}$ at $t = 3096$ days, we checked the gas distribution and velocities at 3196 and 4096 days (Fig. 5.2). However, at these times there appeared to be little fallback. Instead, between times of 9 and 10 years we witness a clear increase in the amount of fallback gas.

We began the *100Sep* simulation with the companion close to the giant's surface. This caused gas to be initially unbound in an asymmetric manner, with the star being quite disrupted in the process. The resulting distribution of fallback gas is complex (see Fig. 5.1, bottom panels). On the other hand, the *218Sep* simulation enters the dynamic inspiral considerably more smoothly than the *100Sep* simulation, giving the primary plenty of time to adjust to the potential of the companion. As a result, the companion causes very smooth waves of density, giving the fallback gas a much more symmetric distribution (see Fig. 5.2, bottom panels).

Using the approximation outlined in Section 5.1, we can calculate an approximate timescale for the formation of a fallback circumbinary disc with fallback material. Fig. 4.6 gives a value for the angular momentum of the bound material of approximately $2.5 \times 10^{52} \text{ g cm}^2 \text{ s}^{-1}$, which can be divided by the average radial distance of the bound mass, $h = 3500 R_{\odot}$, and the mass of the bound envelope, $0.44 M_{\odot}$, at $t = 3300$ days, to give $v_{\theta,h} = 1.1 \times 10^5 \text{ cm s}^{-1}$. Equation 5.5 then can be used to estimate the radius of the circumbinary disc to be $\sim 30 R_{\odot}$. This is explained by the fact that the initial angular momentum is larger in the *218Sep* simulation. From this comparison we conclude preliminarily that to study the formation of circumbinary discs with simulations, we must start with a simulation that contains an approximately correct amount of angular momentum. Using Equation 5.6, the fallback time for this disc is estimated to be ~ 10 years, which is much longer than the duration of our simulation after the fast inspiral phase. Once again we need to be able to run our simulations for longer to study fallback further.

Summary, Conclusions and Future Work

We have carried out simulations of the common envelope interaction using the smoothed particle hydrodynamics code, PHANTOM (Price et al., 2015). In these simulations, the primary star was set up by mapping a 1D stellar profile of density and other quantities into a spherical distribution of gas particles. The core of the star was accreted into a sink particle for computational speed, and the star was then stabilised by damping all velocities for several dynamical times. The compact companion was also represented by a sink particle with a predetermined mass.

These simulations have been benchmarked against one carried out by Passy et al. (2012), showing good agreement. Our simulations, like all those carried out previously, resulted in too small a fraction of unbound mass, and in wider post-dynamic inspiral separations than observed. Nandez et al. (2015) carried out SPH simulations with the inclusion of recombination energy, which produced large unbound mass fractions and lower final separations. It is unclear how the addition of recombination energy can lead to both outcomes, as a greater energy would work to unbind more gas, which generally prevents further inspiral. Those simulations, however, started with a compact giant primary, something that alone will tend to result in more compact post-common envelope binaries. It is not clear why, in their attempt to solve this known problem, Nandez et al. (2015) did not consider giants in the same parameter space as previous simulations.

One of the aims of this project was to determine the effect of gas falling back onto the binary after the infall phase. While it is clear that we see some of this infalling gas, we have been unable to run our simulations for long enough to observe the bulk of the material returning to the centre. This is due to the fact that, at some point in all our simulations, gas begins to be unbound very quickly. However, the time at which this occurs and the speed of the unbinding are both resolution dependent. Higher resolution simulations experienced this sudden unbinding at later times, and the decrease was considerably less extreme. The problem likely resides in the low densities near the sink particles as the simulation progresses, but we have as yet been unable to find a solution. Despite the issue of this unphysical unbinding, we have run the simulations for longer than previous simulations.

We have compared simulations starting at different orbital separations and in doing so shown that ignoring the evolution of the system prior to the fast inspiral phase is not a good approximation. By, simulating the Roche-lobe overflow phase, instead of starting with the companion on the surface of the primary, we show that the final separation and bound mass fraction both depend on the initial separation. In addition, the gas distribution after Roche-lobe overflow is substantially different from the spherical distribution seen in the early phase of simulations that start with the companion on the surface of the giant star, leading to a different rapid infall phase.

Furthermore, we attempted to study the tidal interaction phase immediately preceding Roche-lobe overflow. These simulations could not be completed in the available time. However, the early part of the simulation, before Roche-lobe overflow, showed a great degree of stability, where orbital reduction was at the hand of a pseudo-tide driven by numerical viscosity and the oscillation of the giant envelope due to the introduction of the companion into the computational domain. It is possible that the overall larger angular momentum budget and the longer time for the star to adjust to changes may alter the final outcome of the interaction.

There are many potential avenues for a continuation of the research in this field. The primary path to take is to solve the issue of the artificial unbinding. If we are able to run our simulations for longer, with no artificial unbinding, we will be able to fully examine the behaviour of fallback gas. A second research avenue is investigating the addition of recombination energy. As previously mentioned, Nandez et al. (2015) simulated common envelope interactions using an equation of state that included recombination energy, which resulted in the ejection of the entire envelope. However, it is clear that more work must be done in this area to determine the full range of effects that recombination energy should have on the evolution of a common envelope. The first step would be to determine whether the addition of recombination energy in the simulation by Passy et al. (2012, reproduced here with PHANTOM) helps to eject the envelope and its effect on the final separation. Once established, we would likely need to investigate the details of the energy release to determine the extent to which this energy is available; while our simulations are adiabatic, in nature some of the energy would be radiated away.

Finally, we need to complete our simulations involving the corotating frame. To be able to fully determine the effect of a corotating star on the system, we will need to stabilise the star in the corotating frame before proceeding to the common envelope, something that we have not yet done. This opens the way for a simulation that takes advantage of all of these effects: a simulation including recombination energy, with a corotating primary star placed at the beginning of Roche-lobe overflow (or even farther out) and simulated until any remaining bound gas has fallen back onto the binary. Such a simulation would be far more realistic than previous ones and opens the way for the addition of radiation and magnetic fields, both of which may be of importance to the interaction.



Appendix A: Table of Simulations

| Sim. no. | t_f (yrs) | h_{soft} (R_\odot) | n_{part} | R_1 (R_\odot) | a_i | a_f | a_{1000} | $M_{\text{b},1000}$ (M_\odot) | Cor. | Time (days) | Max Cores |
|-------------|----------------|------------------------------------|-------------------|------------------------|-------|-------|------------|--------------------------------------|------|----------------|--------------|
| 1 | 2.74 | 3 | 76252 | 90.06 | 100 | 15.5 | 16.5 | 0.410 | No | 3 | 32 |
| 2 | 5.28 | 3 | 225164 | 92.52 | 100 | 14.0 | 15.9 | 0.429 | No | 12 | 32 |
| 3 | 5.05 | 3 | 1070079 | 93.05 | 100 | 14.4 | 15.6 | 0.427 | No | 39 | 14 |
| 4 | 5.05 | 3 | 2338260 | 93.97 | 100 | 14.5 | 15.7 | 0.425 | No | 106 | 32 |
| 5 | 10.10 | 3 | 76191 | 88.65 | 218 | 22.4 | 23.4 | 0.330 | No | 10 | 16 |
| 6 | 10.10 | 3 | 225162 | 90.95 | 218 | 24.0 | 22.0 | 0.416 | No | 22 | 14 |
| 7 | 22.37 | 3 | 225162 | 90.95 | 218 | 22.4 | 24.0 | 0.394 | Yes | 55 | 32 |
| 8 | 4.36 | 3 | 109967 | 89.45 | 300 | 299.5 | - | - | No | 30 | 14 |
| 9 | 2.91 | 3 | 225162 | 90.95 | 300 | 299.8 | - | - | No | 77 | 14 |

Table A.1: Parameters for the simulations mentioned in this work; t_f is the length of time for which the simulation was run, h_{soft} is the softening length of the sink particles, n_{part} is the number of particles in the simulation, R_1 is the radius of the primary at $t = 0$, a_i is the initial orbital separation, a_f is the final orbital separation, a_{1000} is the orbital separation 1000 days after the beginning of inspiral and $M_{\text{b},1000}$ is the bound envelope mass 1000 days after the beginning of inspiral.

Table A.1 gives an overview of the basic parameters of the different simulations that were carried out for this work. There are some other simulations that were carried out behind the scene in order to calibrate the code and attempt to solve the artificial unbinding issue, but they were not included in this set. Note also that the simulations labelled 5 and 8 in Table A.1 were not included in the results of the thesis, but are the beginnings of a resolution test for the 218 and 300 R_\odot initial separation simulations.

B

Appendix B: Movies and Code

For visualisations of some of the simulations that have been discussed in the thesis, movies can be viewed by following the links below. The final link gives the Fortran 90 code that has been used for analysis of the simulations.

100 R_{\odot} initial separation, 2.3×10^6 particles, inertial reference frame
<http://web.science.mq.edu.au/~orsola/Thomas/100Movie.mp4>

218 R_{\odot} initial separation, 2.3×10^5 particles, inertial reference frame
<http://web.science.mq.edu.au/~orsola/Thomas/218Movie.mp4>

218 R_{\odot} initial separation, 2.3×10^5 particles, corotating reference frame
<http://web.science.mq.edu.au/~orsola/Thomas/218CorMovie.mp4>

Further, for a movie of Fig. 3.9 (comparison of density around the sinks for different resolutions), please visit:

<http://web.science.mq.edu.au/~orsola/Thomas/SinkEvacuation.mp4>

The code used to calculate energies, angular momenta, orbital separation and bound/unbound mass can be viewed below. Note that this code is dependent on functions defined elsewhere in the PHANTOM code, and hence cannot be run independently.

<http://web.science.mq.edu.au/~orsola/Thomas/FortranCode.pdf>

References

- Bond, H. E.: 2000, in *Asymmetrical Planetary Nebulae II: From Origins to Microstructures*, Vol. 199, p. 115
- Darwin, G. H.: 1879, *Proceedings of the Royal Society of London* **29(196-199)**, 168
- De Marco, O., Passy, J.-C., Moe, M., Herwig, F., Mac Low, M.-M., and Paxton, B.: 2011, *Monthly Notices of the Royal Astronomical Society* **411(4)**, 2277
- Dewi, J. D. and Tauris, T. M.: 2000, *arXiv preprint astro-ph/0007034*
- Djorgovski, S., Drake, A., Mahabal, A., Graham, M., Donalek, C., Williams, R., Beshore, E., Larson, S., Prieto, J., Catelan, M., et al.: 2011, *arXiv preprint arXiv:1102.5004*
- Eggleton, P. P.: 1983, *The Astrophysical Journal* **268**, 368
- Fink, M., Hillebrandt, W., and Röpke, F.: 2007, *Astronomy & Astrophysics* **476(3)**, 1133
- Fryxell, B., Olson, K., Ricker, P., Timmes, F., Zingale, M., Lamb, D., MacNeice, P., Rosner, R., Truran, J., and Tufo, H.: 2000, *The Astrophysical Journal Supplement Series* **131(1)**, 273
- Gafton, E. and Rosswog, S.: 2011, *Monthly Notices of the Royal Astronomical Society* **418(2)**, 770
- Han, Z. and Podsiadlowski, P.: 2004, *Monthly Notices of the Royal Astronomical Society* **350(4)**, 1301
- Harry, G. M., Collaboration, L. S., et al.: 2010, *Classical and Quantum Gravity* **27(8)**, 084006
- Herwig, F.: 2000, *Astronomy and Astrophysics* **360**, 952
- Hurley, J. R., Tout, C. A., and Pols, O. R.: 2002, *Monthly Notices of the Royal Astronomical Society* **329(4)**, 897
- Iben, I. and Livio, M.: 1993, *Publications of the Astronomical Society of the Pacific* pp 1373–1406
- Ivanova, N., Justham, S., Chen, X., De Marco, O., Fryer, C., Gaburov, E., Ge, H., Glebbeek, E., Han, Z., Li, X.-D., et al.: 2013, *The Astronomy and Astrophysics Review* **21(1)**, 1
- Ivanova, N., Justham, S., and Podsiadlowski, P.: 2015, *Monthly Notices of the Royal Astronomical Society* **447(3)**, 2181

- Ivezic, Z., Axelrod, T., Brandt, W., Burke, D., Claver, C., Connolly, A., Cook, K., Gee, P., Gilmore, D., Jacoby, S., et al.: 2008, *Serbian Astronomical Journal* **176**, 1
- Kashi, A. and Soker, N.: 2011, *Monthly Notices of the Royal Astronomical Society* **417**(2), 1466
- Kasliwal, M. M.: 2013, in *IAU Symposium*, Vol. 281, pp 9–16
- Kasliwal, M. M., Kulkarni, S., Gal-Yam, A., Nugent, P. E., Sullivan, M., Bildsten, L., Yaron, O., Perets, H. B., Arcavi, I., Ben-Ami, S., et al.: 2012, *The Astrophysical Journal* **755**(2), 161
- Kuruwita, R., Staff, J., and De Marco, O.: 2015, *Fall-back discs and the end of the common envelope*, In press
- Liu, M., Liu, G., and Lam, K.: 2003, *Journal of Computational and Applied Mathematics* **155**(2), 263
- MacLeod, M. and Ramirez-Ruiz, E.: 2015, *The Astrophysical Journal* **803**(1), 41
- Meyer, F. and Meyer-Hofmeister, E.: 1979, *Astronomy and Astrophysics* **78**, 167
- Miszalski, B., Acker, A., Moffat, A., Parker, Q., and Udalski, A.: 2009, *Astronomy & Astrophysics* **496**(3), 813
- Monaghan, J. J.: 2005, *Reports on progress in physics* **68**(8), 1703
- Morris, J. and Monaghan, J.: 1997, *Journal of Computational Physics* **136**(1), 41
- Mustill, A. J. and Villaver, E.: 2012, *The Astrophysical Journal* **761**(2), 121
- Nandez, J. L., Ivanova, N., and Lombardi, J.: 2015, *Monthly Notices of the Royal Astronomical Society: Letters* **450**(1), L39
- Nandez, J. L., Ivanova, N., and Lombardi Jr, J.: 2014, *The Astrophysical Journal* **786**(1), 39
- Ostriker, E. C.: 1999, *The Astrophysical Journal* **513**(1), 252
- Paczynski, B.: 1971, *Annual Review of Astronomy and Astrophysics* **9**, 183
- Paczynski, B.: 1976, in *Structure and Evolution of Close Binary Systems*, Vol. 73, p. 75
- Passy, J.-C., De Marco, O., Fryer, C. L., Herwig, F., Diehl, S., Oishi, J. S., Mac Low, M.-M., Bryan, G. L., and Rockefeller, G.: 2012, *The Astrophysical Journal* **744**(1), 52
- Price, D., Tricco, T., Wurster, J., Nixon, C., Toupin, S., Laibe, G., Pettitt, A., Dobbs, C., Glover, S., Worpel, H., Nealon, R., Kaluza, O., and Lodato, G.: 2015, *PHANTOM: A smoothed particle hydrodynamics and magnetohydrodynamics code for astrophysics*, In preparation
- Price, D. J.: 2007, *Publications of the Astronomical Society of Australia* **24**(03), 159
- Price, D. J.: 2011, *arXiv preprint arXiv:1111.1259*
- Rasio, F. A. and Livio, M.: 1996, *The Astrophysical Journal* **471**(1), 366

- Rau, A., Kulkarni, S. R., Law, N. M., Bloom, J. S., Ciardi, D., Djorgovski, G. S., Fox, D. B., Gal-Yam, A., Grillmair, C. C., Kasliwal, M. M., et al.: 2009, *Publications of the Astronomical Society of the Pacific* **121**(886), 1334
- Ricker, P. M. and Taam, R. E.: 2008, *The Astrophysical Journal Letters* **672**(1), L41
- Ricker, P. M. and Taam, R. E.: 2012, *The Astrophysical Journal* **746**(1), 74
- Riess, A. G., Filippenko, A. V., Challis, P., Clocchiatti, A., Diercks, A., Garnavich, P. M., Gilliland, R. L., Hogan, C. J., Jha, S., Kirshner, R. P., et al.: 1998, *The Astronomical Journal* **116**(3), 1009
- Sackmann, I., Boothroyd, A. I., Kraemer, K. E., et al.: 1993, *The Astrophysical Journal* **418**, 457
- Sandquist, E. L., Taam, R. E., Chen, X., Bodenheimer, P., and Burkert, A.: 1998, *The Astrophysical Journal* **500**(2), 909
- Schmidt, B. P., Suntzeff, N. B., Phillips, M. M., Schommer, R. A., Clocchiatti, A., Kirshner, R. P., Garnavich, P., Challis, P., Leibundgut, B., Spyromilio, J., et al.: 1998, *The Astrophysical Journal* **507**(1), 46
- Schreiber, M. R. and Gänsicke, B.: 2003, *Astronomy & Astrophysics* **406**(1), 305
- Soker, N. and Kashi, A.: 2011, *arXiv preprint arXiv:1107.3454*
- Staff, J., De Marco, O., Wood, P., Galaviz, P., and Passy, J.-C.: 2015, *Hydrodynamic simulations of the interaction between giant stars and planets*, In press
- Surace, J., Laher, R., Masci, F., Grillmair, C., and Helou, G.: 2015, *arXiv preprint arXiv:1501.06007*
- Tassoul, J.-L. and Tassoul, M.: 1992, *The Astrophysical Journal* **395**, 259
- Tocknell, J., De Marco, O., and Wardle, M.: 2014, *Monthly Notices of the Royal Astronomical Society* **439**(2)
- Tylenda, R., Hajduk, M., Kamiński, T., Udalski, A., Soszyński, I., Szymański, M., Kubiak, M., Pietrzyński, G., Poleski, R., Ulaczyk, K., et al.: 2011, *Astronomy & Astrophysics* **528**, A114
- Webbink, R.: 1984, *The Astrophysical Journal* **277**, 355
- Webbink, R.: 2007, *arXiv preprint arXiv:0704.0280*
- Zahn, J.-P.: 1977, *Astronomy and Astrophysics* **57**, 383
- Zahn, J.-P.: 1989, *Astronomy and Astrophysics* **220**, 112