

THE IMPACT OF INITIAL CONDITIONS ON SIMULATIONS OF THE COMMON ENVELOPE BINARY INTERACTION

By

Roberto Iaconi

A THESIS SUBMITTED TO MACQUARIE UNIVERSITY
FOR THE DEGREE OF DOCTOR OF PHILOSOPHY
DEPARTMENT OF PHYSICS AND ASTRONOMY
NOVEMBER 2016



MACQUARIE
University
SYDNEY • AUSTRALIA

EXAMINER'S COPY

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.

Roberto Iaconi

Acknowledgements

I would like to thank first of all my supervisor, Orsola De Marco, who supported me during the entire time of my Ph.D. by being always available and by helping me with both astrophysical and personal problems. My parents, for letting me make the choices I wanted and helping me in achieving what I wanted. Jan, for helping me with any sort of issue and teaching me a lot of useful things, astrophysical and not.

I would like to thank Caterina, for being there, as best as she could, for any problem I had during these years and for sharing with me some fantastic moments.

A big thanks also to Daniel Price, who let me be a part of his PHANTOM code project and who taught me a lot of very useful numerical concepts.

Thanks to Alessandro, for helping me with all my judo-related injuries, for tolerating all my complaining about them and for his friendship.

Finally, I would like to thank Furio, Evaristo and Sebastiano for all the support and suggestions they always gave me, all the students and postdocs with who I shared my Ph.D. years and everyone at my judo club.

Contributions

We included in this thesis two publications for which not all the work was carried out by me. Additionally, some parts of the work described in this thesis, that have not been published yet, have not been carried out by me. Below we describe in detail my contributions to each publication and to the non-published results that are the outcome of collaborations.

The investigation carried out in Chapter 3, has been submitted for publication to Monthly Notices of the Royal Astronomical Society (MNRAS, <http://mnras.oxfordjournals.org>) on 02/March/2016. I am the lead author of the submitted paper and have carried out all of the tasks described therein except for the two SPH simulations carried out as comparisons, that were performed by Master student Thomas Reichardt, although I had originally set up a stable star within the SPH code PHANTOM, which had never been used to perform CE simulations before.

The set of five simulations analysed in Chapter 5 were performed by our collaborator J.-C. Passy using Canadian supercomputing resources. The analysis of the outputs of these simulations was entirely carried out by me.

The study carried out in Chapter 6 was carried out after the discovery of an energy conservation issue with the ENZO code as applied to our problem. This issue was unveiled by me during the analysis of the simulation data for the work described in Chapter 3. Due to the investigation carried out in Chapter 6, which resulted in a viable solution to the problem that was affecting other simulations that were being carried out within our group, I became co-author of the paper by

Staff et al. (2016a).

Abstract

In this work we analyse various aspects of the common envelope (CE) interaction between two stars via numerical simulations.

The common envelope (CE) interaction is a short phase of the interaction between two stars (a primary and a companion) in a binary system characterised by the dense cores of the two objects orbiting inside their merged envelopes. During this phase, orbital energy and angular momentum are transferred to the gas of the envelope, that can become unbound from the potential well of the system, leaving behind a close binary. Unfortunately, due to its short duration, the CE phase is not readily observed (only one case has been observed until now) and numerical simulations are a major way to investigate its physics. However, to this time, numerical studies have failed to fully reproduce the observed post-CE parameters, that is, close binaries with separations generally shorter than $\simeq 5 R_{\odot}$ and where all the envelope has been expelled, yielding instead rather large final separations and never expelling the whole envelope.

Since the CE interaction has been analysed in multiple works, but without explicitly taking in consideration the effect of the single physical parameters, in this PhD we tried to do so.

One of the main topics we investigated during this PhD work has been the effect of large initial binary separations on the CE interaction. We performed a simulation with the binary components initially placed at the maximum possible distance that would guarantee the system to end in a CE. The main outcomes of this work show that a larger initial separation does not dramatically affect the CE interaction. The final separations obtained in this way are slightly larger with respect to an identical system where only the separation is reduced in such a way that the CE begins at the beginning of the simulation. The amount of mass unbound from the potential well of the binary is

also slightly larger.

Another important part of this work has been the study of the effects of rotation on the CE interaction. To achieve this goal we spun up the original star we used for the study on larger separations, after investigating the possibility and reliability to create a more accurate stellar model. The results of this investigation show that initial rotation of the primary star has negligible effects on the outputs of the CE interaction.

The third effect we worked on is the variation of the final separation and unbound mass in function of the mass of the primary. We therefore performed a set of simulations with a more massive primary star and set of companions with different masses. This simple study showed that for the same companion's mass a more massive primary generates a closer binary at the end of the CE interaction, in the range of observations, yielding however less unbound mass.

Additionally, during the work we encountered a numerical problem with ENZO, which showed poor conservation of energy in our simulations. We therefore had to devote part of this PhD work to investigate the issue and find a solution for it.

Contents

Acknowledgements	v
Contributions	vii
Abstract	ix
1 Introduction	1
1.1 Stellar evolution	1
1.2 Binary interactions	3
1.3 The common envelope interaction	7
1.3.1 The hypothesis of the common envelope phase	7
1.3.2 The importance of studying common envelope evolution	9
1.4 Numerical simulations of common envelope binaries	11
1.5 Motivation for the current project and outline	14
1.5.1 Simulating a CE binary system with a rotating giant star	14
1.5.2 The impact of the pre rapid in-spiral phase on CE simulations outcomes . .	15
1.5.3 The impact of a massive primary star on the CE output	15
1.5.4 Code-code comparison	15
2 Our “codes” and the generation of the initial stellar model	17
2.1 The three modelling codes used in this thesis	17
2.1.1 The hydrodynamic code ENZO	18

2.1.2	The hydrodynamic code PHANTOM	21
2.1.3	The stellar structure and evolution code MESA	23
2.2	Common envelope simulations: initial stellar modelling	25
2.2.1	Creation of the one-dimensional stellar model	25
2.2.2	Mapping the 1D stellar model into the 3D computational domain	26
2.2.3	Criteria to evaluate damping and relaxation	33
2.2.4	A stellar relaxation simulation	34
3	The effect of the initial separation on common envelope simulations	41
3.1	Introduction	42
3.2	ENZO simulation setup	43
3.2.1	Single star setup and stabilisation	44
3.2.2	Binary system setup	45
3.3	Results	47
3.3.1	Orbital Separation	47
3.3.2	Envelope ejection	51
3.3.3	Tidal bulges	58
3.3.4	Evolution of the gas velocities and density in proximity to the companion: the end of the in-spi	
3.3.5	Angular momentum and energy conservation	63
3.4	Comparison with PHANTOM simulations	65
3.5	Comparison with published simulations	68
3.5.1	Side-by-side code comparison	69
3.5.2	The final orbital separation as a function of M_2/M_1	69
3.5.3	The final orbital separation as a function of primary mass or envelope binding energy	70
3.5.4	The final orbital separation as a function of giant spin at the time of Roche-lobe overflow	70
3.5.5	Unbound mass	71
3.6	Conclusions	72
4	The effect of the primary rotation on the common envelope simulations	75
4.1	The rotation of red giant branch stars	76
4.2	One dimensional simulations with rotation	77

4.2.1	Generalities of the simulations	78
4.2.2	Varying the ZAMS angular frequency	79
4.2.3	Varying η_* and M_*	82
4.2.4	A work-around of code limitations for fast rotating stars	83
4.3	Three dimensional simulations of the CE phase with rotation	87
4.3.1	Simulation setup	87
4.3.2	Evolution of the binary separation	89
4.3.3	Envelope ejection	93
5	The effect of massive primaries on the common envelope interaction	99
5.1	Simulations setup	100
5.2	Evolution of the binary separations	104
5.3	Envelope ejection	108
5.4	Conclusions	111
6	Energy conservation in ENZO	115
6.1	The problem of energy conservation in grid codes	115
6.2	Evaluating energy conservation	117
6.2.1	Calculation of energies	118
6.3	Estimating the extent of energy non-conservation	119
6.4	Further tests	121
6.5	Solving the energy conservation issue	125
7	Conclusions and future work	127
7.1	An unsolved physical problem	127
7.1.1	On envelope recombination energy simulations	128
7.2	The aim of this project	129
7.3	Results	130
7.3.1	CE simulations with a rotating giant	130
7.3.2	CE simulations with a wide initial orbital separation	131
7.3.3	CE simulations with more bound primaries	132

7.3.4	Energy conservation in common envelope simulations	133
7.4	Future work	134
A	Appendix A: co-authored publications	137
	References	139

1

Introduction

To contextualise the work made in this thesis it is important to introduce basic literature and concepts involving the evolution of single and binary stars, followed by a more detailed description of the main phenomenon that will be studied throughout this thesis: the common envelope binary interaction.

1.1 Stellar evolution

Stars form out of clouds of interstellar matter. Some of these clouds can collapse and form a proto-star, that is basically a dense clump of gas contracting on a thermal time-scale. The molecular hydrogen making up the original cloud gas dissociates and ionises as the temperature rises. Eventually nuclear fusion is triggered (at $\sim 10^7$ K, this condition can be fulfilled only for proto-stars with mass $\gtrsim 0.08 M_{\odot}$).

When a star begins its life, it is said to enter an evolutionary phase called main sequence (MS). During the MS the luminosity of the star is produced by hydrogen nuclear burning in the core, while the rest of the star is basically inactive and radiatively or convectively expelling the energy produced. Hydrogen burning produces as an output helium, which starts accumulating in the core but is not yet burning due to the low temperatures. The core may be either convective or radiative depending on the mass of the star.

When central hydrogen burning stops, the star has a helium core surrounded by a hydrogen-rich envelope burning at the base. The hydrogen burning shell temperature increases due to the heating created by the contraction of the inert stellar core and the star expands and, eventually, grows in luminosity. We call these first-time giants, or red giant branch (RGB) stars. In this phase, due to the increased radius and luminosity the star is very prone to lose mass via stellar winds. Eventually core contraction triggers core helium burning and the star leaves the RGB phase. It is during the RGB phase that the star has its first opportunity to interact with a nearby stellar or planetary companion.

Helium burning in the core terminates when the material is processed principally into carbon and oxygen and, similarly to what happened before when hydrogen stopped burning in the core, helium continues to be burned in a shell surrounding the exhausted core. Now the star has two shell sources: the more internal helium shell and the hydrogen burning shell above it. While the helium shell burns outwards the CO core increases in mass and contracts. At the same time hydrogen burning dumps helium on the shell, which eventually ignites (quite violently in an event called a helium shell flash).

During helium and hydrogen shell burning the star grows in radius and luminosity for the second time, in a phase called the asymptotic giant branch (AGB). Similarly to what happens during the RGB phase, because of the large radius and intense luminosity, the star tends to lose mass via stellar winds. During the last phases of the AGB the mass-loss rate intensifies until all but $10^{-2} - 10^{-3} M_{\odot}$ of envelope material is left, at which point the remaining envelope contracts and the stellar effective temperature is seen to increase. The material expelled by the winds can be heated and ionised by the radiation from the central object; as a result of this ionisation the shell of gas emits and can be observed as a planetary nebula (PN). Eventually the nuclear shell sources run low on fuel and are quenched. The core is now a white dwarf (WD).

For the more massive stars (mass greater than $8 M_{\odot}$) nuclear burning of the carbon in the core can take place. At increasing stellar mass, heavier elements can be synthesised and burned (in the order: oxygen, neon, magnesium, silicon, sulphur, iron and nickel), forming as a consequence a more complex onion-like structure. After these subsequent burning phases massive stars end their life exploding as core-collapse supernovae, ejecting their entire envelope and leaving behind a neutron star or black hole remnant. WDs are supported by the pressure of highly degenerated electrons and have radii and densities of the order of, respectively, $\simeq 10^{-2} R_{\odot}$ and 10^6 g cm^{-3} . Neutron stars are instead supported by the pressure of partially degenerate neutrons and have radii and densities of the order of $\simeq 10^{-5} R_{\odot}$ (10 km) and $10^{14} \text{ g cm}^{-3}$, respectively.

For a more detailed description of the various stages of the stellar evolution, we refer the reader to Kippenhahn et al. (2012).

1.2 Binary interactions

A non-negligible percentage of the stars in the universe are known to be part of multiple systems, mostly binaries. As a result of these statistics many classes of objects and types of physical phenomena are known to originate from binary stars or are hypothesized to be so. Additionally, stars must also often interact with their planets, since it is known that planets are often present around main sequence stars. The discovery of planets close to evolved stars plausibly indicates that some kind of interaction has taken place.

The binary fraction, that is the percentage of binary systems among the total number of systems, where a system is intended as a single star or a multiple star system, changes as a function of the mass of the primary ($\lesssim 0.1 M_{\odot}$ = very low mass, $\simeq 0.1 - 0.5 M_{\odot}$ = low mass, $\simeq 0.7 - 1.3 M_{\odot}$ = solar type, $\simeq 1.5 - 5 M_{\odot}$ = intermediate mass, $> 8 M_{\odot}$ = massive; Duchêne & Kraus 2013). More massive stars are more prone to having a companion, showing a binary fraction $> 60\%$. Instead intermediate mass, solar type and low / very low mass stars have binary fractions respectively $> 50\%$, 41% and 22% . The binary fractions seem to be independent of whether the stars are in clusters or in the field. Together with the binary fraction, the period, the ellipticity and the mass ratio distribution determine the frequency and intensity of the binary interactions that can take place, driving the evolution of systems. The mass ratio distribution is approximately flat for solar

type primaries in wide orbits, while in close binaries similar mass components are somewhat more probable. Intermediate mass and massive stars seem instead more likely to exist in non-equal mass systems for both wide and close binaries. There is, however, a great deal of uncertainty in the observations (Duchêne & Kraus 2013).

Most of the binary systems have large separations and the components are not likely to have a substantial impact on each other. However, if the stars are close enough, i.e., with orbital separations of at most a few stellar radii of the larger star, the components can directly interact exchanging matter in different ways. Common mechanisms mass transfer via Roche-lobe overflow and wind mass transfer. In this work we will focus on the former. When the components of a binary system are close enough they tidally deform and spin-orbit interaction circularises the orbits. If this configuration remains stable the system is defined as a detached binary. However, at some point, instabilities, such as the Darwin instability (Darwin 1879), may develop. Darwin instability takes place when one of the components of a stable binary system grows in size because of stellar evolution. The components of a stable binary system orbit in co-rotation, that is, the orbital frequency and the spin frequencies of the stars around their own axes are the same. This yields the lowest energy configuration for the binary. When one of the stars increases in size it spins down, due to conservation of angular momentum, as a result the orbit tries to reach a new minimum energy configuration where the system is stable by giving angular momentum to the star. This spins the star up, slows the orbit down and reduces the orbital separation, re-establishing co-rotation. However, there is a threshold (Darwin instability) over which the orbit cannot give enough angular momentum to the star to re-establish co-rotation, roughly when the spin angular momentum of the star is greater than one third of the orbital angular momentum, keeping the orbit from becoming stable again and pulling the stellar components closer. This generates a new tidal phase that brings the components close enough that mass transfer can take place. Mass transfer is a complicated process that involves a lot of physical mechanisms, but is basically driven by the combined effective potential field of the two stars, where “effective” means that all the fictitious forces, such as the centrifugal force, are included in the potential. The combined potential of the two objects forms a field of equipotential surfaces around them with special points of equilibrium called Lagrangian points (Figure 1.1). If one of the stars fills the equipotential surface corresponding to the Lagrangian point L_1 , part of the envelope material can escape through the very same point to




FIGURE 1.1: General schematic of equipotential surfaces and Lagrangian points L_{1-5} around stars 1 and 2 with masses m_1 and m_2 , respectively. If we suppose that Star 2 is the accreting component filling its Roche lobe, L_2 and L_3 are respectively the outer Lagrangian point on the secondary and primary side. Image from Benacquista & Downing (2013).

be drawn towards the companion star. This mass transfer can begin for example if the primary expands beyond the above-mentioned equipotential surface or if the orbital size is decreased by tides. When mass transfer is ongoing through L_1 the system is defined a semi-detached binary. Due to the orbital motion, mass does not accrete directly onto the companion, but generates an accretion disk around the companion. In some cases, for example in the case of mass transfer with a small mass ratio, both stars can fill the L_1 equipotential surface. The two stars share in this case a common atmosphere which fills the equipotential surface going through L_2 , and are defined a contact binary.

From the moment when mass transfer begins the type of behaviour we can expect depends sensitively on the interplay of a large number of stellar and orbital parameters. Mass transfer can be conservative (no loss of angular momentum from the system) or non-conservative, stable or catastrophic. Binaries can have one or more episodes of mass transfer as the stars in the systems change as they evolve or as a result of the interaction itself. It goes without saying that the phenomenology is complex and extremely difficult to interpret. There is no need to list the details of each process, but it is interesting to describe some common classes of objects, observed or theorised, that are a by-product of a binary interaction (see Carroll & Ostlie 2006 and references therein for additional

details):

1. Algols: main sequence or sub-giant stars simply undergoing mass transfer. Semi-detached binary systems where the originally more massive star has transferred enough mass to its companion to invert the mass ratio.
2. Cataclysmic variables and nova-like binaries: in these systems the components are very close and always involve a white dwarf and usually a main sequence star filling the Roche lobe (semi-detached binary). Examples are dwarf novae and classical novae, binaries that release energy in a very powerful event related to the accretion onto the white dwarf. The peculiarity of these binaries is that the binary survives the release of energy and the outburst event can happen again. On the other hand Type Ia supernovae exhibit a more energetic outburst, again based on accretion onto a white dwarf that increases its mass above the Chandrasekhar limit, leading to a thermonuclear runaway that destroys the binary.
3. X-ray binaries with neutron star and black hole components: systems characterised by a bright emission of X-rays and comprising, in general, a neutron star or, in a smaller number of cases, a black hole. These objects are powered by accretion from a non-degenerate companion.
4. Symbiotic binaries: relatively long period interacting binaries, composed of a giant primary and a white dwarf. The giant's stellar wind accretes onto the compact companion. Occasionally the giant can be filling its Roche lobe.
5. Barium and S-star binaries: here there is no ongoing binary interaction. They are long period binaries composed of a giant companion and a white dwarf primary with an observed abundance rich in elements that the giant cannot have manufactured itself. It is widely believed that the white dwarf donated those elements to the giant by an interaction some time in the past.
6. Post-common-envelope objects: of particular interest since specifically related to the topic of this thesis, they include a wide number of binary types. The majority of the systems described in the previous points may also be included in this category, since CE could, in

certain cases, be a part of their evolution. However, we will highlight here those binaries for which CE is the last step of evolution prior to their current state. These objects are in most part short period binaries composed of a hot dwarf or sub-dwarf and a main sequence star or other compact objects. Presumably the former was originally a giant or sub-giant stripped of its envelope as a result of a CE interaction. Planetary nebulae with a close binary in their centre likely derive from CE interactions. Mergers are another aspect of this category; some CE events could lead to close binaries that then dissipated the remaining orbital energy via gravitational waves and ended up forming a single star. Some mergers could happen inside the CE (Tylenda et al. 2011)

1.3 The common envelope interaction

To understand the evolution of stellar binary systems it is essential to analyse the interactions between their stellar components. Such interactions include physical phenomena such as spin-orbit tidal interactions or mass transfer. All of these processes can act together to give rise to complex binary interaction phenomena such as outflows and outbursts, for example in novae and some supernovae. In this thesis we study one of these possible binary interactions, known as the common envelope (CE) interaction. It occurs when two stars in a close binary spiral-in towards each other within a single extended atmosphere, or envelope, after one of the two expands.

Since it has been first hypothesized by Paczynski (1976), its general features have been approximately understood. However, it is the details of the physics behind the CE interaction that dictate the terms of the engagement and what type of binary is forged by this process. Current CE models cannot explain the observations of post-CE binaries. For a comprehensive review of the CE interaction we point the reader to the following reviews: Iben & Livio (1993), Taam & Sandquist (2000), Webbink (2008), Taam & Ricker (2010) and Ivanova et al. (2013).

1.3.1 The hypothesis of the common envelope phase

The idea of the CE interaction dates back to the 1970s, when the discovery of the binary object V471 Tau in the Hyades cluster was identified as a possible cataclysmic binary system. Because

of its relative proximity to us, it was relatively easy to measure its parameters with good precision. From these parameters it was determined that this binary must have gone through a CE phase.

Cataclysmic binary systems such as novae and dwarf novae were known to be short period binaries consisting of a WD primary and a main sequence companion donating mass to the primary via an accretion disk. From stellar evolution theory the expectation for the masses of the components is around $1 M_{\odot}$ for both the stars, and that the WD star was originally grown as the core of a red giant or super-giant star. The red giant would have been larger than today's orbital separation, implying that the companion must have been inside the giant in the past. That means that the two stars must have gone through a CE phase.

When the giant progenitor of the WD expanded and filled its Roche lobe, a contact system was formed, and soon thereafter the common surface of contact could expand beyond the outer Lagrangian points (see Figure 1.1) without being lost immediately. This can be analytically shown if the envelope is not solid-body rotating (and there is no obvious reason for the envelope of a binary system to rotate as a rigid body), hence theory suggests that a binary of this type could be a reasonably long-lived object. The transition between a contact binary and a CE binary is defined as the points at which the common surface expands beyond the outer Lagrangian points. With the components of the binary orbiting within the CE (these components are the degenerate core of the giant and the companion star), there will be a gravitational drag force on the stars, that makes them lose orbital energy and in-spiral towards one another. With this basic physical understanding, Paczynski (1976) proposed that:

$$\frac{\tau_D}{P_{\text{orb}}} \sim \frac{M}{a^3 \rho_{\text{CE}}}, \quad (1.1)$$

where τ_D is the time-scale on which the drag force makes the binary spiral-in, P_{orb} is the orbital period, $M = m_1 + m_2$ is the total mass of the system, a is the orbital separation and ρ_{CE} is the average density of the CE. This gives a criterion to estimate the spiral-in time as a function of the ratio between the average density of the matter within the separation a , given by $\rho_a \sim M/a^3$, and ρ_{CE} . Equation 1.1 is useful to estimate the time-scale of a CE interaction provided observational data of a system are available. In addition, various important considerations were made by Paczynski (1976):

- Through the drag force, the orbital energy and the angular momentum of the binary system

are passed to the CE gas, respectively heating it and spinning it up.

- As the in-spiral goes on, it will become more difficult to transfer energy from the orbit to the CE. Intuitively, the relative velocity of the components with respect to the surrounding matter will decrease, because the envelope material will be close to co-rotation, decreasing in turn the drag force.
- Although the in-spiral slows down (see previous point), it could inject enough energy in the CE to eject it and leave the binary system as we observe it now, while the ejected matter, if ionized by the heat of the degenerate star, could be observed as a planetary nebula.

Finally, it is worth noting that the estimation of τ_D for the system considered in that (V471 Tau) paper was $\tau_D \simeq 10^3 P_{\text{orb}}$, and since $P_{\text{orb}} \simeq 12 \text{ hours} = 0.5 \text{ days}$, we get a in-spiral time-scale of $\tau_D \simeq 500 \text{ days} \simeq 1.4 \text{ years}$. This time is of the same order of magnitude of the dynamical times of giants (months to years).

With these considerations, Paczynski (1976) laid the foundations of the CE interaction, but left many details unresolved. Unfortunately, without a detailed description we cannot even start to make predictions of what type of binary systems emerge from CE interactions.

1.3.2 The importance of studying common envelope evolution

The CE interaction is the gateway to the formation of a rich set of astronomical objects, such as cataclysmic variables, gamma ray bursts and Type Ia supernova progenitors.

Once a CE phase begins, the ejection of the envelope is probably not inevitable. In general, we expect that if the envelope is ejected a close binary emerges from the interaction. However, if the CE is not ejected the most likely fate for the central stars is a merger. Today there is still no quantitative understanding of what leads to one or the other outcome.

Various astrophysical objects have gone through one or more CE phases as introduced in Section 1.2 (see Ivanova et al. 2013 and references therein):

- High-mass compact binary mergers. The merging of compact objects in systems composed of a double neutron star, a double black hole or a neutron star plus a black hole, is thought to produce detectable gravitational waves. The recent detection of gravitational waves was

in fact consistent with a black hole-black hole merger (Abbott et al. 2016a, Abbott et al. 2016b). These kind of systems are expected to have gone through one or two CE events. Therefore, understanding how to create a close binary system that can later evolve in one of the merger types mentioned above could help interpret future observations of gravitational waves by constraining their merger rates in stellar population synthesis models.

- Low-mass compact binary mergers. Another type of compact merger is potentially responsible for at least some type Ia supernovae, thought to be the merger of two carbon-oxygen white dwarfs. This scenario is similar to the one above: a system with a wide orbit has to shrink and expel the gaseous envelope, then the remaining close binary system will merge through other physical processes. Also in this case understanding the physics of the CE would help constrain the type Ia supernova rate through stellar population synthesis, interpret the observations and establish the elusive nature of the progenitors.
- Type Ia supernovae from single degenerate systems. Binaries composed of a main sequence or a giant star transferring mass onto a carbon-oxygen white dwarf are believed to be another channel for the production of type Ia supernovae. Most of these binaries went through a CE phase so an improved understanding of the CE phase could help to understand their production.
- Gamma-ray bursts. The short duration gamma-ray bursts have been hypothesized to be the result of compact mergers, hence they could be the results of the phenomena introduced above. The long duration gamma-ray bursts type should derive instead from particular types of core-collapse in massive stars that require the loss of the envelope, achievable through the CE. As in the previous cases, we could make more accurate predictions of their rates and characteristics with an improved understanding of the CE phase.
- The physics of CE interactions also has the potential to give us a different way to look at the evolution of single stars. A star might have its evolution altered by CE-type in-spiral of a planet, a brown-dwarf or a low-mass main sequence star. These interactions could have observable consequences such as polluting the stellar envelope, spinning it up or even ejecting it (Passy et al. 2012, Staff et al. 2016b). The ejection due to the planets could in turn

explain the formation of single low-mass white dwarfs.

The ejected envelope could evolve into a planetary nebula if the effective temperature of one or both of the central stars were enough to ionize the gas (as originally stated by Paczynski 1976). Observations of planetary nebulae show that at least about 10-15% of their population has a close binary system in the centre (Bond 2000, Miszalski et al. 2009a). Planetary nebulae form when stars eject the envelope at the end of the AGB via dust-driven winds and/or binary interactions. Similarly, the formation via a CE phase sees the CE ejection mechanism as the way to expel the envelope of the AGB star.

Moreover, a CE ejection naturally explains the axi-symmetric shapes of planetary nebulae typically observed around post-CE binary central stars (Miszalski et al. 2009b). In fact the great majority of planetary nebulae are not spherical, but they display axi-symmetries and point symmetries, sometimes with the addition of jet-like structures. In general these kind of shapes have been explained through particular models of stellar rotation and magnetic fields (García-Segura et al. 1999), but there are counter arguments that show that in the majority of the cases these physical phenomena cannot be sustained if the star is single: the presence of a companion could drive the ejection through preferential ways (e.g. equatorial mass loss) and sustain a magnetic field to create jets (Nordhaus & Blackman 2006, Nordhaus et al. 2007; but see also De Marco 2009 for a complete discussion of the problem). The physical mechanisms producing ejecta and jets as a result of a CE interaction had already been analysed by Soker (1997) and updated later by De Marco & Soker (2011), who distinguish five types of planetary nebula shaping binary interaction.

To conclude, the importance of studying the CE resides in the fact that there are many unsolved questions around close binary systems both from the theoretical and observational points of view, and the understanding of the physics of the phenomenon could give both the initial conditions to model more clearly all of its “children” and a lot of useful prediction for the observations.

1.4 Numerical simulations of common envelope binaries

To discuss the general approach used to study the CE; it is important to note that the phenomenon is so complex that a complete analytical model is impossible, and all analytical estimates can only be used as guidance. The CE interaction is also inherently 3D and 1D numerical models are

not accurate enough to describe it. However, so far research that needs prescriptions for the CE phase, such as population synthesis calculations, have used the analytical approach. For example modeling CE interactions in stellar population synthesis studies has relied upon simple analytical prescriptions based on energetic or angular momentum considerations (e.g. Politano & Weiler 2007, Nelemans 2010, De Marco et al. 2011). These back-of-the-envelope analytical estimates have been essential to direct the approach to the study of CE with numerical simulations.

Current 3D CE hydrodynamic simulations have at their basis a simplified model for the binary system, due to the fact that this type of hydrodynamical simulations requires substantial allocation of computational time, using expensive supercomputing facilities. Several numerical studies have been undertaken. Primary examples are those of Rasio & Livio (1996), Sandquist et al. (1998) and Sandquist et al. (2000) for older references, while to cite more recent important work we have Ricker & Taam (2008), P12: Passy et al. (2012) (hereafter P12), Ricker & Taam (2012), Nandez et al. (2014), Nandez et al. (2015), Ohlmann et al. (2016), Staff et al. (2016a) and Staff et al. (2016b). All of them include similar ingredients but test different types of numerical codes (see in particular P12 for the comparison of different code methods). The complexity of modern 3D hydrodynamic codes gives rise to difficulties when interpreting the outputs of simulations. Examples of physical processes investigated using simulations are: the efficiency of hydrodynamical and gravitational drag forces, accretion onto the companion, the morphology of the outflowing CE, the effects of gas recombination on the energy budget and many others. However, the big question, which still has to be answered, is how the CE is ejected, and how the post-CE binary parameters relate to the input parameters.

There are several reasons why it is so difficult to achieve the envelope ejection in numerical simulations. The CE evolution can be divided into different phases, with different time-scales, lengthscales and different driving physical mechanisms. Following Podsiadlowski (2001) and Taam & Ricker (2010) the following phases characterise the whole phenomenon:

1. During the initial phase the two stars, still detached and with no mass exchange, approach one another due to tidal torques. If co-rotation between the orbit and the stellar rotation is reached, orbital decay may still take place because of the Darwin instability. When the primary fills its Roche lobe, mass transfer between the primary giant and the smaller companion

starts.

2. The second phase is identified as a rapid in-spiral. After an initial short period of orbital decay and mass exchange through Roche-lobe overflow, the rapid in-spiral takes place. In this phase energy and angular momentum are deposited in the CE, driving its expansion and may lead to its dynamical ejection. As noted before, the time-scale of this phase is very short, around one year.
3. The final phase may include a slow in-spiral phase. Once the orbital decay of the system in the CE has deposited enough energy in it to cause significant expansion, the rapid in-spiral will slow down (assuming the CE has not been ejected yet) to a stage where all the frictional energy that is released is transported to the surface where it is radiated away, until the envelope is eventually completely ejected by the percentage of energy that is not lost by radiation. This phase can last hundreds of years.

As one can see, the three phases have very different time-scales and are driven by different physical processes, hence it is simply impossible, at the moment, to reproduce all of them in a single numerical simulation. What is done instead are numerical simulations of the single phases independently, but this does not account for the natural continuity of the phenomenon. Additionally, the number of physical processes involved in the CE phase is huge and each of them is a free parameter (sometimes constrained by observations) that has unpredictable consequences. Not all these parameters have been explored yet. For example, factors like the initial rotation profiles of the components, recombination energy of the CE gas during the expansion and the effect of layers of CE material not properly ejected that fall back onto the binary, have only been investigated partially.

While simulating the entire interaction will not be possible in the immediate future, the most achievable way to investigate CE is to systematically add new physical parameters to present simulations and determine the changes in the outputs with respect to previously obtained results.

1.5 Motivation for the current project and outline

This thesis' project aims to gauge the impact of the input parameters on the outcomes of 3D hydrodynamic simulations. Previous numerical work investigated primarily individual simulations with a particular set of initial conditions. However, without any comparison simulations, carried out with different initial parameters, one cannot properly gauge the impact of different physical mechanisms on the outcome of the CE. Notable exceptions are the work of Sandquist et al. (1998) and of P12, where such investigations did take place and which we will discuss fully throughout this thesis. Here we investigated in detail three parameters, which can be chosen as initial conditions: the rotation of the primary star, the initial orbital separation, the mass of the primary's envelope.

Finally, we also tested CE simulations using a range of computational techniques: the hydrodynamic code ENZO in uniform, static grid as well as adaptive mesh refinement (AMR) modes and of the smooth particle hydrodynamics (SPH) code PHANTOM, used here for CE simulations for the first time.

1.5.1 Simulating a CE binary system with a rotating giant star

This sub-project focuses on adding the rotation of the primary star to one of the CE simulations of P12. The rotation of the primary star has been included by Rasio & Livio (1996), Sandquist et al. (1998), Ricker & Taam (2012) and Ohlmann et al. (2016). All but Rasio & Livio (1996) use grid codes in the inertial frame, with the primary rotation velocity set as a function of the orbital velocity to be co-rotating or slightly below co-rotation. Rasio & Livio (1996) instead set their binary up in the co-rotating frame. The former papers do not stabilise the giants in the co-rotating frame, while the latter paper does. None of the published papers takes into account realistic rotational profiles for giant stars. Sandquist et al. (1998) are the only ones who compare two rotating and non-rotating giant models with otherwise identical parameters. They concluded that rotation only affects marginally the ejection of the envelope. The other publications do not discuss the effects of rotation on the outcome of the simulation.

The questions that we set out to answer therefore are: (i) does the lack of giant stabilisation impact the CE outcome? (ii) Is it possible to set up a giant with a realistic rotation profile? (iii) Would rotation affect the outcome of the CE simulation carried out by P12?

1.5.2 The impact of the pre rapid in-spiral phase on CE simulations outcomes

For this simulation we set the companion orbiting at a larger distance from the primary than in the simulations of P12, but close enough that a tidal capture is expected. With a bigger simulation box compared to that of P12, so as to contain a system with a greater separation, we will also be able to follow the ejecta further. Although simulating tidal capture correctly is beyond our computational capabilities, angular momentum conservation insures that all the orbital angular momentum is carried into the CE. The pre in-spiral phase comprises stellar spin up as well as Roche-lobe overflow. Both these phases likely alter the energetic exchange prior to the in-spiral compared to a simulation where the companion started on the surface of the primary.

1.5.3 The impact of a massive primary star on the CE output

It is likely that a more massive and/or more compact primary will impact the CE outcome. A given post-CE primary mass can be created by a more massive star caught into a CE earlier during its RGB evolution, or a less massive one caught later. At that time, a more massive star will be more compact than a lower mass one for equal core mass. Being more compact and more massive could be an advantage to reach a smaller final separation and consequently extract more orbital energy and possibly also eject more mass. We therefore set out to compare a simulation with a larger RGB envelope mass to the simulation of P12.

1.5.4 Code-code comparison

Previous numerical work on the CE interaction have used both grid (e.g., Sandquist et al. 1998, Ricker & Taam 2012), SPH (e.g., Rasio & Livio 1996, Nandez et al. 2015) and more recently moving-mesh (Ohlmann et al. 2016) codes, but only P12 actually compared different numerical techniques (grid and SPH) using the same setup. We follow the same approach, comparing the results of the grid code ENZO and of the SPH code PHANTOM (see Chapter 2 for a description of the software). This gives more robustness to our results and additionally gives us the possibility to use PHANTOM for the CE problem for the first time.

2

Our “codes” and the generation of the initial stellar model

2.1 The three modelling codes used in this thesis

Modern astrophysics widely uses numerical simulations to reproduce any kind of interaction. As a result a myriad of different codes are available. It is therefore generally possible to simulate the same astrophysical problem with multiple codes. CE evolution itself had been simulated before with different of them. Since the major part of this thesis work involved modelling of single stars and binary systems with numerical tools, we had the necessity to chose the codes to fit our purposes. In this chapter we describe the three codes we used to carry out all of our work: ENZO, PHANTOM, and MESA.

2.1.1 The hydrodynamic code ENZO

Before describing the ENZO code (O’Shea et al. 2004, Bryan et al. 2014) in detail, it is useful to show the hydrodynamic equations that are solved by hydrodynamic codes that includes self-gravity. These equations can be used in various forms depending on the formalism of the code (e.g., Eulerian vs. Lagrangian). However, here we will only show the Eulerian form used to simulate fluids in ENZO. ENZO can solve the above-mentioned equations with or without magnetic fields, cosmological expansion (the code was originally developed for cosmological simulations), radiative cooling, radiative heating and conduction, but for our purpose these ingredients will be omitted. As a consequence the equations we are solving take the following form:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{v}) = 0 , \quad (2.1)$$

$$\frac{\partial \rho \vec{v}}{\partial t} + \rho(\nabla \cdot \vec{v})\vec{v} + \rho(\vec{v} \cdot \nabla)\vec{v} = -\nabla p - \rho \nabla \Phi , \quad (2.2)$$

$$\frac{\partial e}{\partial t} + \nabla \cdot [(e + p)\vec{v}] = -\rho \vec{v} \cdot \nabla \Phi , \quad (2.3)$$

where ρ is the density, t is time, \vec{v} is the velocity, p is the pressure, Φ is the gravitational potential and e is the total energy per unit volume ($e = u + (\rho v^2)/2$), where u is thermal energy per unit volume). The first equation represents conservation of mass, the second of momentum and the third of total energy. The system is closed by an ideal gas equation of state

$$p = u(\gamma - 1) , \quad (2.4)$$

where γ is the ratio of specific heats, and by the Poisson’s equation for the gravitational potential

$$\nabla^2 \Phi = 4\pi G \rho , \quad (2.5)$$

where G is the gravitational constant. This set of equations is numerically solved at each time-step for each element of fluid simulated. In the following paragraphs we briefly describe how the simulated fluid is mapped and how the hydrodynamic equations are solved.

ENZO is a block structured adaptive mesh refinement (AMR), parallel code meant to provide both high spatial and temporal resolution for modelling astrophysical fluid flows. The code is Eulerian, can be run in all the three dimensions (1D, 2D, 3D) in a Cartesian grid and supports a wide variety of physics: hydrodynamics, ideal and non-ideal MHD, N-body (self-gravity of fluids and particles), primordial gas chemistry, optically thin radiative cooling of primordial and metal enriched plasmas (as well as some optically thick cooling models), radiation transport, cosmological expansion and models for star formation and feedback in a cosmological context.

The AMR algorithm of ENZO uses an adaptive hierarchy of grid patches at varying levels of resolution. Each of these patches covers a region in space requiring higher resolution than its parent grid and can itself become a parent grid to a higher refinement level. This hierarchy is reconstructed constantly during the simulation since it has to adapt to the emergence of “interesting” regions in the grid. AMR can be extended to arbitrary depths by adding additional grids to refine regions in the sub-grids. AMR allows different methods of refinement to generate sub-grids from a parent grid. One of the most commonly used methods is to refine by gas mass, on the assumption that regions of denser gas are “interesting”. A cell is refined if the gas mass contained in it is greater than a certain value. This method mimics the Lagrangian method by trying to keep a fixed mass resolution. AMR presented certain grid boundary problems when tested on CE simulations until a new solver was designed by Passy & Bryan (2014). With the new solver energy is well conserved for the point-mass particles representing the core of the primary and the companion for CE interactions (see Section 2.2, Chapter 3 and Chapter 6). In the work presented here ENZO was used in static grid mode, i.e., using a uniformly spaced grid set up at the beginning of the simulations as well as in AMR mode, once the AMR solver became available.

ENZO allows various hydrodynamic methods to solve the fluid equations. What we used for all our simulations is the second-order finite difference finite difference scheme previously used in the ZEUS code (Stone & Norman 1992a). It is important to note that ENZO is independent from ZEUS, only the latter algorithm has been implemented. This algorithm is peculiar because uses a staggered mesh such that velocities are face-centred, while density and internal energy are cell-centred. This algorithm is formally second-order accurate in space and first-order accurate in time. It is an aperture-split method, which breaks the solution of the fluid equations into parts that are solved successively. Equations are split into three parts: operator-split expansion terms (i.e.,

the terms dependent on the derivative of the acceleration), source terms and transport terms (i.e., advection terms for mass, momentum and energy). Self-gravity is instead handled by using a fast Fourier technique to solve the Poisson equation on the root grid at each time-step.

The code, in addition to gas, can include collisionless matter, modelled with particles that interact with other particles and gas only via gravity. Particles are not adaptively refined, instead each particle is associated with the highest refined level available at its position in the domain and are moved between grids as the grid is rebuilt between one time-step and the following. As a result, a particle has the same time-step and feels the same gravitational force as a grid cell at that refinement level. To calculate the dynamics of collisionless systems, ENZO uses a particle-mesh N-body method and to avoid non-physical point mass effects, particles are considered point masses only until a certain level of refinement. Above this threshold, level contributions from particles to the gravitating mass field will be smoothed over a spherical region centred at the particle position.

All grids on a level are advanced with the same time-step. The time-step at a certain level is computed by first calculating the largest time-step allowed for each cell, then the level is advanced with a time-step equal to the minimum of all of the values previously computed. Since the code is grid based, it has a spatial limit given by the simulation box. Matter that reaches the limits of this box is handled by a range of boundary conditions. For our problem we used “vacuum” boundary conditions (see Section 2.2), such that outflowing gas leaves the domain and cannot return nor interact further.

ENZO outputs can be saved to disk with any required cadence, but these dumps can be very memory expensive, especially when simulating at high resolution or with AMR. Therefore one usually chooses a dumping frequency fitting the time-scale of the simulated problem, but not too high to avoid excessive memory usage and slowing down of the simulation due to the input-output times. For example, in CE simulations, the frequency at which outputs are typically saved is every 0.01 to 0.05 yr, on total simulation times of the order of a couple of years. However, ENZO also gives the possibility of having higher time cadence outputs for data analysis done *during* the course of a simulation. The cadence of this data analysis can be as high as every sub-cycle of the finest refinement level, without the need to write an entire dump. This is done through the YT data analysis package (Turk et al. 2011). During the computation ENZO can be linked to YT so

that the mesh geometry, fluid quantities and particle data are saved as *NumPy*¹ arrays. YT then processes these data in the same way it would read data from a data dump saved on the disk. This processing is done via a user-defined *Python* script, which is run at a user-defined frequency during the simulation. Every time the script is run, ENZO does not proceed to the next time-step until the data analysis is completed.

ENZO is parallelised using Message Passing Interface (MPI). Parallelisation is accomplished by distributing grids among processors, where the root grid is split into a number of tiles typically equal to the number of processors. Sub-grids are then moved to the least busy processors to balance the computational load.

2.1.2 The hydrodynamic code PHANTOM

A second hydrodynamic code used for this thesis is the smoothed particle hydrodynamics (SPH) code PHANTOM. PHANTOM (Price & Federrath 2010, Lodato & Price 2010) is a code designed to avoid high memory usage and its original purpose was primarily to model star formation.

The need to have a comparison code arises from the fact that grid and SPH codes have their distinctive advantages and disadvantages, therefore using both of them on the same physical problem allows one to estimate with more precision possible shortfalls of the numerical method. The major advantages of grid-codes with respect to SPH are for example better resolution of shocks, better treatment of dynamical instabilities (e.g. Kelvin-Helmholtz, Rayleigh-Taylor instabilities) and better resolution of low density regions. On the other hand, advantages of the SPH codes over grid-codes are their intrinsic adaptivity, without the need for additional numerical methods like AMR, the capability to follow all the mass, no advection issues and Galilean invariancy. This has advantages when checking on energy conservation.

A code comparison between grid-codes and SPH on CE simulations has been carried out by P12, using ENZO (see Section 2.1.1) and the SPH code SNSPH (Fryer et al. 2006). The results of their investigation show that on a “standard” CE setup, namely a giant primary star with a companion orbiting at its surface in circular orbit, the two codes are in very good agreement. However, the comparison has not been extended to different regimes. Additionally, the SPH code

¹<http://www.numpy.org/>

PHANTOM was never used for CE simulations before. Therefore we also analyse the CE evolution and its outputs with a new code (see Chapter 3).

PHANTOM solves the fluid equations in Lagrangian form, i.e., the fluid is discretised into a set of “particles” that move with the fluid velocity field. The properties of one of these particles are determined as a function of the neighbouring ones, therefore what are usually called particles are in fact interpolation points.

The density field is computed as:

$$\rho_a = \sum_b m_b W(|\vec{r}_a - \vec{r}_b|, h_a) , \quad (2.6)$$

where a and b are particle indexes, ρ_a is the the density at particle a and m_b is the mass of particle b . W is the smoothing kernel, r_a and r_b are the locations of particles a and b with respect to a chosen centre of reference and h_a is the “smoothing length” of the considered particle. The sum is over neighbouring particles. The smoothing length for particle a is given by:

$$h_a = h_{\text{fact}} \left(\frac{m_a}{\rho_a} \right)^{\frac{1}{3}} , \quad (2.7)$$

where h_{fact} is a proportionality factor specifying the smoothing length in terms of the mean local particle spacing. This formula only works for equal mass particles, a condition enforced in PHANTOM. Since ρ_a and h_a have a mutual dependence, the two equations have to be solved simultaneously. The kernel function is instead given by:

$$W = \frac{C_{\text{norm}}}{h_a^3} f(q) , \quad (2.8)$$

where C_{norm} is a normalisation constant (π^{-1} in three dimensions) and $f(q)$ is a cubic spline function with $q = |\vec{r}_a - \vec{r}_b|/h_a$.

Neighbour finding is done via a kd-tree as a default and its implementation closely follows that of Gafton & Rosswog (2011), splitting particles recursively based on the centre of mass and bisecting the longest axis at each level. The tree is refined until a cell contains less than N_{min} particles. N_{min} is set by default to 10. At each of these nodes the code then executes neighbour finding. After this procedure is performed the density field is calculated.

Domain decomposition is used in PHANTOM to apply a shared memory parallelisation (OpenMP).

PHANTOM has options to solve the gas equations in different forms. In all of these forms the equation of motion for each particle is discretised and numerically solved. The default form, also used in this work, is simple compressive hydrodynamics plus an adiabatic equation of state, using $\gamma = 5/3$. In addition to this basic gas equation, PHANTOM can solve magneto-hydrodynamics in both the ideal and non-ideal cases (dissipative magnetic field), dust-gas mixtures, chemistry with cooling and particles injection.

The time-stepping, while integrating the equations of motion, is done by a leap-frog integrator. This particular version of the integrator is both reversible and symplectic, making the algorithm Hamiltonian: linear and angular momentum are exactly conserved, while energy and phase-space volume are very well conserved.

The time-step is limited by the Courant condition and it is chosen to be the minimum between a set of constrained physical quantities from the simulation (accelerations, external forces, sink-particle accelerations, cooling function). The time-step can also be individual for each particle for problems with a large range of time-scales over the domain. This increases the efficiency (in terms of time) of the code, increasing the speed by approximately one order of magnitude, but reducing the conservation of physical quantities.

Sink particles are available in PHANTOM. This type of particle is different from SPH particles in that they interact with other particles (other sink particles and SPH particles) only via gravity. For sink particles, gravitational interaction is computed by direct summation; they can accrete gas and can store all the physical quantities of the accreted SPH particles. In this work we have used sink particles as fixed mass particles, or accreting only for a short amount of time at the very beginning of the simulation, to generate a particle with the mass of the giant star's core (see Chapter 3).

2.1.3 The stellar structure and evolution code MESA

The code we used to model the giant primaries prior to inserting them into the 3D hydrodynamic codes is MESA and a detailed description of the code can be found in Paxton et al. (2010),

Paxton et al. (2011), Paxton et al. (2013) and Paxton et al. (2015). Here we give a general description of the features of the code.

Solving the coupled, non-linear differential equations of stellar structure has to be done computationally. MESA represents one of the most recent codes devoted to the solution of stellar structure equations and is an open-source code that combines many numerical and physics modules for 1D stellar simulations. Additional modules can be easily added to increase the scope of the code. In addition to the core part of the software, to model stellar structure and evolution, MESA is designed to be used for a wide range of stellar physics studies. Concepts such as equation of state, nuclear reactions or opacities can be applied to a wider range of phenomena. By default a wide range of stellar evolutionary scenarios can be simulated: from low- to high-mass stars and from pre-main sequence to advanced evolutionary phases.

The code is OpenMP parallel and allows for both AMR and adaptive time-stepping to optimise the level of detail of the physics where needed and to improve performance. The way MESA operates is to build a 1D (spherically symmetric) model where the number of cells depends on the complexity of nuclear burning, gradient of state variables, composition, etc. The equations of stellar structure and composition, physically coupled, are solved simultaneously numerically and include sophisticated equations of state, opacity, nuclear reaction rates, element diffusion and atmosphere boundary conditions. Below we will discuss some of these in more detail.

A detailed zone-by-zone equation of state incurs a certain computational burden. To alleviate this burden and speed up the simulations it is convenient to compute tables before the simulation starts, which are then read and used during the simulation. Similarly, thermonuclear and weak reactions are tabulated before the beginning of the computation. The nuclear reactions are instead computed during the runs and various reaction networks are available.

In addition to micro-physics MESA handles numerous macro-physical phenomena happening during stellar evolution, including various mixing length treatments for convection, convective overshooting, treatment of the stellar atmosphere outside the photosphere (to guarantee stability), diffusion, gravitational settling and different possible wind schemes. Additionally, MESA also implements schemes to simulate binary interactions; for example it is possible to calculate two individual stars that are transferring mass between one another.

The code has been verified (i.e., the physical equations are correctly solved) and validated (i.e.,

it is solving the correct set of physical equations) against various other codes and observations for a wide range of problems. In addition the code is constantly being developed, thanks to a large community of users.

2.2 Common envelope simulations: initial stellar modelling

The main purpose of this work is to simulate the interaction between an RGB star and a compact companion, paying particular attention to details such as the impact of the initial structure of the giant primary on the simulation outcomes. Here we review the numerical procedure used, giving technical references where needed.

In this work we expand on the work of P12, by evaluating a greater number of quantities and details that may have repercussions on the interaction simulations. The initial stellar models used in our 3D simulations are pre-calculated in 1D using MESA and they are physically very accurate. However, after it is transposed into the hydrodynamic codes becomes unphysical: it has no energy source, it does not radiate and it is governed by a simplistic equation of state. It is, however, in hydrostatic equilibrium, and has a reasonably accurate density and internal energy profile. This makes it sufficiently realistic in simulations that take place on the dynamical time-scale. Here we explain some aspects of this setup.

2.2.1 Creation of the one-dimensional stellar model

For some of our simulations we used the same stellar structure calculation as P12, where a $1 M_{\odot}$ star was evolved until its hydrogen-exhausted core mass reached $M_{\text{core}} = 0.392 M_{\odot}$, with the EVOL code (Herwig 2000). At this point the star has a total radius of $R = 83 R_{\odot}$ and a total mass $M = 0.88 M_{\odot}$, smaller than the initial value because of mass loss. These older stellar models were used for consistency with P12. However, we carried out equivalent MESA models, which confirm that they are almost identical and that they can therefore be used interchangeably.

2.2.2 Mapping the 1D stellar model into the 3D computational domain

In Section 2.2.1 we have described how we use one-dimensional stellar evolution codes to generate a stellar structure model of an RGB star. The next step is to insert the stellar model into the ENZO computational domain and stabilise it. In Figure 2.1 we show how density, temperature and pressure of the 1D model result once mapped into the 3D code ENZO. The most important things




FIGURE 2.1: Comparison between the 1D model (blue) used for the simulations in Chapters 3 and 4 and the same model once mapped in 3D (red). Panel (a) is density, (b) is temperature and (c) is pressure.

to highlight are the difference in resolution between the two codes and the jump in temperature on the outside of the star after the mapping in the 3D domain. The reasons of these features are described below.

One of the main differences between MESA and ENZO is the equation of state used by the two codes. While MESA uses a tabulated equation of state that takes into account factors such

as different elements, radiation, recombination and ionisation, specific to reproduce the physics taking place inside a star during its evolution, ENZO uses a more generic ideal equation of state for mono-atomic gas (i.e., the entire envelope is made of mono-atomic hydrogen). As a result not all the MESA equation of state related quantities need to be imported into ENZO. For each radius value of the 1D model we only import density and pressure, from these quantities the initial temperature can be extrapolated through the ideal gas equation of state. After the quantities are interpolated to the 3D domain the first time, the hydrodynamical solver of ENZO evolves them and no additional inputs are needed from the 1D model. The central energy source of the star, that is the core, where nuclear reactions take place, it is not reproduced in ENZO. However, the pressure imported from the 1D model is the result of the MESA calculation, that includes nuclear core burning. With these pressure values we are able to produce a stable stellar model. Being the equilibrium of the star well reproduced, the lack of the central energy source is not expected to affect the CE interaction from an hydrodynamical point of view. The radiation produced by the core could affect the CE by applying radiation pressure to the envelope material. This aspect of the interaction has not yet been studied and we leave it to future work.

The stellar relaxation procedure that we are about to explain is reasonably standard and was first adopted by Rasio & Livio (1996). On the other hand, publications using this procedure lack a quantitative and detailed explanation of this process. As simulations gain in sophistication so will the importance of accurate model mapping and good stability. We here quantify our procedure and the degree to which it achieves stabilisation so as to provide a reference point for future work. As one can see from Figure 2.2, a stellar model imported from a 1D stellar evolution code evolved in ENZO without stabilisation loses the original structure and starts expanding in the simulation domain within one dynamical time (21 days for our model). If one could follow the model with a bigger simulation domain and for its whole natural relaxation time, eventually the matter would stabilize to a new spherical structure around the central point-mass. However, we are limited by both domain sizes and simulation time. Therefore the necessity to obtain artificially a stable model arises. Below we describe the procedure to map the 1D star into the computational domain, followed by the stabilisation procedure. While several of the parameters are the same as for the standard ENZO code, others have been created especially for the CE problem by P12. We list here those parameters that are particularly important, justifying their values.




FIGURE 2.2: Density slices on the $x-y$ plane at $z=0$ for our RGB stellar model at a resolution of 128^3 cells and with a computational domain of 2 AU per side, without velocity damping. The plots show different times: 0 yr (a), 0.29 yr (b) and 0.58 yr (c). Arrows show the velocity field. The small green triangle shows the position of the central point-mass particle.

Other parameters have their default values, which can be found in the ENZO user guide²:

- *ProblemType* = 41; enables the modifications carried out by P12 to simulate binary common envelope interactions.
- *TopGridDimensions* = 256 256 256; grid resolution.

²<http://enzo-project.org/doc/index.html>

- *LeftFaceBoundaryCondition / RightFaceBoundaryCondition* = 5 5 5; vacuum boundary conditions are implemented for this problem. The classical outflow boundary conditions created a significant artificial mass in-fall into the simulation domain as mass is attracted into the domain from the ghost regions surrounding it by the mass in the box.
- *LengthUnits* = 6e13; 4 AU in centimetres. The value of this parameter represents also the simulation box size if it is not specified through any other parameter. The box size is double the one of P12. We took as a reference the 128^3 cells simulations in that paper, and to keep the star resolved with the same number of cells our linear resolution was doubled to 256^3 (see the parameter above).
- *TimeUnits* = 3.1558e7; 1 yr in seconds. The time-scale of a CE interaction is of the order of years, hence the year is the natural time unit.
- *MassUnits* = 1.9891e33; 1 M_{\odot} in grams.
- *BinaryStarHaloDensityFract* = 1e-4; this parameter lets the user chose the density of the gas surrounding the star as a fraction of the density of the surface shell of the star. The halo gas parameters other than density are set by the code such that $P_{\text{halo}} = P_{\text{outer stellar shell}}$ (where P is the pressure) to avoid stellar instabilities due to the missing of a realistic stellar atmosphere which would gradually vanish in the surrounding medium. Since ENZO uses an ideal gas equation of state, the very low pressure gas surrounding the star is at very high temperature. This is evident from Figure 2.1 (b).
- *BinaryStarUsePointMasses* = 1; 0 = standard ENZO dark matter particles, 1 = point masses. P12 modified the potential of a dimensionless particle to be defined analytically so as to reduce the uncertainty in gravity computations. For details see P12, section 2.3, paragraph 5.
- *BinaryStarNumberOfParticles* = 1; 1 = single star, 2 = with companion. Here we are running a simulation to stabilize our giant star in ENZO, before running the CE simulation, hence we insert just one body.

- *BinaryStarCentralMass* = 0.3959; in M_{\odot} . Counter-intuitively, this value is not the mass of the core of the giant star. To explain clearly what this parameter represents, a description of how ENZO interpolates the 1D stellar model to a less refined 3D model is given below.
- *BinaryStarFile* = *RGB_SPH_Init.txt*; name of the 1D stellar structure model file.
- *StaticHierarchy* = 1; 0: AMR, 1: no AMR. Before late 2014 we used ENZO in uniform, static grid mode. When upgrades by Passy & Bryan (2014) became available we started using the AMR capability.
- *ParticleCourantSafetyNumber* = 0.4; maximum fraction of a cell width that a particle is allowed to travel per time-step.
- *CourantSafetyNumber* = 0.4; this is the maximum fraction of the CourantFriedrichsLewy-
implied time-step that will be used to advance any grid.
- *PointMassPotentialType* = 1; 1 = Sandquist et al. (1998), 2 = Plummer. This parameter represents the shape of the smoothed potential of the point-mass particles. The Sandquist et al. (1998) potential is smoothed by the prescription described in Equation 6.6.
- *PointMassSmoothingLength* = 3.516e11; in cm. The gravity smoothing length for the point-mass particles: typical values are between 1.5 and 3 times the cell width in uniform, static grid mode and the smaller cell width in AMR. This parameter takes its value in physical units, not in code units, so one has to consider the single cell physical length as a function of the *LengthUnits* parameter, of the grid dimensions and of the box size. The particular value is hence 1.5 times the width of our static main grid in centimetres. The value of the smoothing length is of great consequence to the energy conservation properties of the code and we will discuss the issue further in Chapter 6.
- *RelaxationOn* = 1; 0 = off, 1 = on. Activates damped relaxation of the star. At every time-step the velocities are divided by a factor defined by the next parameter. We use this to speed up the stellar structure relaxation time.
- *RelaxationFactor* = 3.0; when *RelaxationOn* = 1, at every time-step all velocities are divided by this number.

Once parameters are set the 1D star is mapped into ENZO in the following way. To start, ENZO loops over all its 3D cells and computes their centres. Then for every ENZO cell centre its distance, r , from the centre of the domain, ENZO applies a bisection method to the 1D star shell radii, with the first interval being the whole radius (r_0), as schematically shown in Figure 2.3. In this way the shell of the 1D star that corresponds to a given cell centre distance is determined; this means that the value of r is between the lower and the upper bounds of the 1D shell selected with the bisection method.




FIGURE 2.3: Classic bisection method. Here R is the current ENZO shell centre radius, which we called r in the text, and r_0 , r_1 and r_2 are the 1D model radii.

Next, to assign the values of the various physical quantities to the ENZO cell of radius r , we perform a first-order interpolation by averaging their values at the shell's boundaries, respectively r_1 and r_2 , weighing those values by a factor that depends on the distance of r from r_1 and r_2 . That is:

$$\begin{cases} (r - r_1)/(r_2 - r_1) = k, \\ (r_2 - r)/(r_2 - r_1) = 1 - k. \end{cases} \quad (2.9)$$

Then to average the physical quantities, for example the density, ρ , at the centre of the ENZO cell of radius r we have:

$$\rho(r) = \rho_2(1 - k) + \rho_1 k, \quad (2.10)$$

where ρ_1 and ρ_2 are, respectively, the densities at the bottom and top boundaries of the 1D shell. The problem with this method is that the stellar mass, particularly for giants, is very concentrated in the core (see Figure 2.4). Because of the steep density gradient near the core, a large part of the mass is lost in the central ENZO cell during the interpolation. The fact that the entire central part of the star is missing from the 3D model mapped in ENZO it is also clear from Figure 2.1, where the red lines, representing the ENZO model, reach a lower radius of at most $2 R_\odot$. The mass missing because of interpolation must be added back into ENZO to properly represent the

star and preserve equilibrium, and since it is mostly lost in the core, this can be done by adding a point-mass particle of the residual mass at the centre of the star, by choosing an appropriate value for the *BinaryStarCentralMass* parameter. To choose the appropriate value, we estimate the missing mass, by comparing the total mass of the 1D model to the total stellar mass enclosed in the ENZO computational domain after the interpolation. The contribution of the medium surrounding the star can be minimised by choosing an appropriate low density (e.g., *BinaryStarHaloDensityFract* = 10^{-20}).




FIGURE 2.4: A representation of how ENZO samples a 1D model stellar density profile. A generic stellar density profile trend is plotted. Evident is the very steep gradient near the core, where most of the mass of the star is enclosed.

The procedure carried to set up a stellar model in the SPH code PHANTOM is mostly analogous to what was discussed above for ENZO. The first difference is that, during the process of interpolating the data, PHANTOM generates an user-defined number of SPH particle with equal mass. The sum all the particles' masses corresponds to the total mass of the star. Then all the SPH particles in a certain radius around the centre (decided in function of the simulation) are removed and replaced by a point-mass with a total mass equal to that of the particles removed. If this was not done, the extreme density of SPH particles in the core would dramatically slow down the simulations. The final model obtained in this way has therefore the same structure of the ENZO one. The second difference is that in SPH no external low-pressure, high-density gas is required, since the simulation domain is only limited to the space occupied by the existing particles.

Furthermore, the damping and relaxation procedure described in the following sections and the evaluation criteria are identical for both ENZO and PHANTOM, therefore we limit ourselves to their description in the former.

2.2.3 Criteria to evaluate damping and relaxation

The 1D stellar evolution codes are extremely refined both from the point of view of the physics and of the grid resolution. ENZO's grid is coarser and its adiabatic equation of state is different from MESA. As soon as the star is mapped into ENZO it starts to expand. We speed up the relaxation process by dividing all velocities on the grid by a factor three at every time-step.

To determine when the star is relaxed we compare the gas velocities to typical orbital velocity values, because the entire interaction takes of the order of one orbital period. We therefore do not want velocities due to instabilities to be larger than the physical motion of the gas we are trying to study. The criterion is therefore:

$$v_{\text{num,max}} < f v_{\text{Kep}} , \quad (2.11)$$

where $v_{\text{num,max}}$ is the maximum of all gas velocities at the selected simulation output, v_{Kep} is the Keplerian circular velocity of the companion orbiting at the stellar surface or

$$v_{\text{Kep}} = \sqrt{\frac{G(M+m)}{a}} \simeq 33.7 \text{ km s}^{-1} , \quad (2.12)$$

for a $0.6 M_{\odot}$ companion orbiting at 3 stellar radii ($= 249 R_{\odot}$) and f is a tolerance factor, which is not chosen *a priori*. These parameters can be adapted to the problem.

We also compare our gas velocities with the local sound speed ($C_s = \sqrt{\gamma P / \rho}$, where $\gamma = 5/3$, P is pressure and ρ is density). This comparison gives an estimate of the stability of the model on a local scale. In fact if in a certain region the gas velocities are bigger than C_s , unwanted shocks could propagate through the model.

We run the damping for more than one dynamical time:

$$t_{\text{dyn}} \simeq \sqrt{\frac{R^3}{2GM}} , \quad (2.13)$$

where R is the radius of the primary star, M is its mass and G is the gravitational constant. At this point the perturbations due to the stabilization had the time to propagate across the star. So, monitoring the velocities for at least $1t_{\text{dyn}}$ gives us an idea of the magnitude of f in the type of stellar models we are using.

2.2.4 A stellar relaxation simulation

We run a relaxation simulation using the ENZO parameter file described in Section 2.1 and the same stellar evolution code file output of P12, to create a primary star for a CE interaction simulation. We run the damping part of the simulation for an arbitrary number of dynamical times ($\simeq 7t_{\text{dyn}} \simeq 0.4$ yr). The evolution of the density vs. radius profile is shown in Figure 2.5 (left panel). The density profiles plotted in Figure 2.5 are averaged over a sphere of fixed radius. Note that in this case we do not overplot the MESA profiles since the initial model (blue line) it is the same as the red line showed in Figure 2.1 (a). The loss of spherical symmetry once the star is mapped in ENZO is minimal, hence the averaged density profiles are almost identical to those taken along any direction. The density profile has a very small variation over the simulation time except in the central part of the star, within $\simeq 10 R_{\odot}$ (recall that the stellar radius is $\simeq 83 R_{\odot}$). The density decreases rapidly in that region at the beginning of the simulation before stabilising. The density profile is steeper and poorly sampled in this region, thus the star needs to reconfigure itself more near the core than at the surface. At the beginning of the simulation velocities are all zero due to the fact that they are not imported from MESA (Figure 2.5, right panel; solid blue line). However, the velocity profile starts to change rapidly as soon as the simulation is started, but the damping is effective in keeping the values of the velocity from growing. At the end of this simulation ($t \simeq 7t_{\text{dyn}}$) velocities are below $\simeq 2 \times 10^{-4} \text{ km s}^{-1}$ (Figure 2.5, right panel; dashed red line). We can now consider the criterion in Equation 2.11 to estimate the tolerance factor f at $\simeq 7t_{\text{dyn}}$. The value of $v_{\text{num,max}}$ is of the order of $2 \times 10^{-4} \text{ km s}^{-1}$, hence

$$f = \frac{v_{\text{num,max}}}{v_{\text{Kep}}} \simeq 6 \times 10^{-6}. \quad (2.14)$$

Since, as we will see below, the stellar model results in reasonable equilibrium if allowed to evolve without damping, we decided that a factor $f = 6 \times 10^{-6}$ is sufficient to damp this type of primary




FIGURE 2.5: Left panel: the averaged radial density as a function of the radius at the beginning (solid blue) and at the end (dashed red) of the damped simulation. Right panel: same as the left panel, but for the shell-averaged radial velocity as a function of radius.

star. This value has then be used as a reference for additional damping simulations carried out in this thesis work. The value of C_s is always far greater than the value of v_{num} in every cell of the box for our stellar model at the end of the damping simulation (see Figure 2.6).




FIGURE 2.6: Values of C_s (blue line) and those of v_{num} (red line) as a function of the linearized cell number (cell have been linearised starting from the the corner of the box with coordinates (0,0,0) and moving along y , then z , then x . Note that not all the cells are plotted, but just those enclosed in a sphere of $100 R_\odot$ radius and centred at the centre of the box. The sphere contains the whole star.

Following the damped simulation we reset all velocities to zero and restart the simulation without velocity damping for a few more dynamical times to ensure the stability of the star. We reset the velocities to zero by accessing the ENZO output files, which are in the HDF5 format, using the python programming language module `h5py`³, then we restart our ENZO simulation from this modified output, changing the settings of the following parameters in the ENZO parameter file to:

- *StopTime* = 0.80931448857348; $14t_{\text{dyn}}$ in code units.
- *RelaxationOn* = 0; relaxation turned off.

A time sequence of slices in the three directions of the Cartesian axes is shown in Figure 2.7 and will be used together with the radial velocities and densities vs. radius profiles for the analysis. In this case the profile shows a more substantial variation at the surface of the star than at the core, where four density cusps form in the directions of the coordinate axes. These cusps are the result of the method used to map the 1D model into the 3D domain, that leaves small singularities at every 90 degrees angle aligned with one of the three directions of the ENZO domain. These singularities evolve in the cusps when the model is relaxed into the ENZO domain. At increasing resolution (either by increasing the overall resolution in uniform grid mode or by using AMR) the magnitude of these extrusions diminishes. We evaluated their mass by considering the density profile along a direction that does not overlap the cusps, for example the 45° direction in the plane $x - y$ at $z = 0$. We selected as radius the point where the density drops to an almost constant value, which results in a sphere of radius $100 R_\odot$, centred at the centre of the domain, as shown in Figure 2.8. The mass enclosed in the sphere is $\simeq 0.87869 M_\odot$, while the total mass including the cusps is $\simeq 0.87874 M_\odot$, hence the mass in the cusps due to numerical effects is $\simeq 5 \times 10^{-5} M_\odot$. In conclusion, the amount of mass which accumulates in the cusps is very small and we expect it to have negligible effects on the binary interaction simulation.

More significant is the radius of the relaxed star as compared to the original MESA radius, as is clear from Figure 2.9 (left panel). The radius has increased by $\simeq 17 R_\odot$ (i.e., from $83 R_\odot$ to $100 R_\odot$) with respect to the 1D stellar radius, but the amount of mass that filled the expanded volume is $\simeq 2 \times 10^{-3} M_\odot$, a very small amount compared to the total mass of the star. So its

³a good explanation on the structure of the ENZO outputs HDF5 file can be found in the official documentation at: http://enzo.readthedocs.org/en/latest/user_guide/HierarchyFile.html; while everything about `h5py` can be found at: www.h5py.org




FIGURE 2.7: Slices on the x (left row), y (central row) and z (right row) axes at times $\simeq 0.40$ yr (first column), $\simeq 0.52$ yr (second column), $\simeq 69$ yr (third column), $\simeq 0.81$ yr (fourth column). The first slice of every column represents the last output of the damping part of the simulation, with all the velocities reset to zero.

effect on the binary interaction should be relatively small. The evolution of the radial velocity vs. radius profile is instead shown in Figure 2.9 (right panel). As soon as the simulation is started without damping velocities develop in the model. Since the star is now readjusting, velocities are chaotic inside the model. Bigger variations appear near the surface at the beginning, where the star is slightly expanding. This behaviour reflects the changes that occur to the density profile. Over the $7t_{\text{dyn}}$ of the un-damped simulation, velocities tend to decrease on the surface and the oscillations propagate towards the core. At the end of the simulation the overall values of velocity remain below $\simeq 3 \text{ km s}^{-1}$. Also in this case to ensure the reliability of the model we applied the criterion in Equation 2.11, using the value of 3 km s^{-1} as a reference for the maximum velocity. Hence:

$$f = \frac{v_{\text{num,max}}}{v_{\text{Kep}}} \simeq 0.1, \quad (2.15)$$




FIGURE 2.8: Zoom of the last panel of the right column of Figure 2.7. The color map has been changed to enhance the cusps and the velocity arrows have been removed. The over-plotted circle has a radius of $100 R_{\odot}$.

and it is clear that in this case the velocities that develop in the star are not completely negligible. However, we have carried out damping and relaxation simulations for longer times without a noticeable increase in the final internal velocities. Additionally, increasing the resolution, e.g., using AMR, does reduce the velocities leftover after stabilisation.

The values of C_s and v_{num} for each cell are shown in Figure 2.10, from which we estimated that the simulation is free of artificial shocks. Finally, we checked the stability of the central point-mass particle between the beginning of the damping simulation and the end of the relaxation simulation. Due to the symmetry of the numerical processes that develop in the model the central point-mass particle does not move during the simulation.

In conclusion, our relaxed stellar model retains its initial spherical symmetry, as well as the giant density profile and it does not exhibit extreme numerical features. However the spurious velocities that are still present after the relaxation are not as low as one might hope. In the future, for more realistic simulations, which may be run for time-scales longer than a few dynamical times we will have to further consider stellar stability. For the time being, this is the best model we can achieve, so we progress to our binary, CE simulations.




FIGURE 2.9: Left panel: the averaged radial density as a function of the radius at the beginning (solid blue) and at the end (dashed red) of the un-damped simulation. Right panel: same as the left panel, but for the shell-averaged radial velocity as a function of radius.




FIGURE 2.10: The same as Figure 2.6 but for the un-damped part of the simulation. In this case the radius of the sphere enclosing the analysed cells is $150 R_{\odot}$.

3

The effect of the initial separation on common envelope simulations

Here we study the outcome of CE simulations where the initial orbital separation is larger than one stellar radius. Several simulations from previous work have initial conditions whereby the companion is close enough to the primary that the latter is overflowing its Roche lobe. We speculate that letting the primary evolve to Roche lobe contact can alter the dynamics compared to a situation where the two stars are initially in contact.

Most of this chapter was presented as a paper. The paper has been submitted to the Monthly Notices of the Royal Astronomical Society (MNRAS, <http://mnras.oxfordjournals.org>) on 02/March/2016. We have edited out information that was already presented or that will be presented in the following chapters. The full submitted paper is available at <http://arxiv.org/abs/1603.01953>. I am the lead author of the submitted paper and

have carried out all of the tasks described therein except for the two SPH simulations carried out as comparisons, that were performed by Master student Thomas Reichardt, although I had originally set up a stable star within the SPH code PHANTOM, which had never been used to perform CE simulations before.

3.1 Introduction

Here we analyse the effect of the initial orbital separation on the final outcome of CE simulations by carrying out a set of simulations that parallel one of the simulations carried out by P12, where a $0.88 M_{\odot}$ red giant branch (RGB) star interacts with its $0.6 M_{\odot}$ compact companion. In their simulation the companion was initially placed near the surface of the giant. In one of our simulations we place instead the companion at the approximate largest distance from which an orbiting companion is likely to be brought into Roche lobe contact with a giant. It is expected that prior to the start of the CE in-spiral phase, tidal forces will redistribute orbital energy and angular momentum from the orbit to the primary. Eventually the primary would overflow its Roche lobe and start mass transfer onto the companion, resulting in the rapid CE in-spiral. These phases are expected to induce envelope rotation and expansion, changing the overall distribution of the envelope and lowering its binding energy. The envelope would be lighter and easier to unbind, but the overall strength of the gravitational drag (Ostriker 1999) may be smaller because of relatively lower densities and smaller velocity contrasts. It is therefore not clear *a priori* what effect a larger initial separation would have on the simulation.

The effect of rotation alone is the central investigation in Chapter 4. Here we hope to obtain rotation as a side effect of the tidal forces. It is however important to highlight that the effect of a rotating giant is difficult to determine from previous work, as only Sandquist et al. (1998) carry out a rotating/non-rotating comparison. Hence, as a part of this chapter, we will also consider the rotation we were able to induce in the primary. Considering instead the initial orbital separation, all simulations from previous work started at a separation such that the giant was already overflowing its Roche lobe and thus could not determine the effects of a more gradual expansion of the giant envelope.

We carry out our simulations with grid (in AMR mode) and SPH codes. In so doing we

compare different numerical techniques while making the most of what each has to offer. The SPH code we use, PHANTOM (Price & Federrath 2010; Lodato & Price 2010), has never been used for CE interaction simulations before, hence this work serves also to introduce PHANTOM to this problem.

3.2 ENZO simulation setup

The simulations of P12 were run with a static uniform grid. ENZO already had AMR capabilities when P12 performed their simulations, but they were not available for CE simulations. However, given the most recent updates applied to ENZO (Passy & Bryan 2014) we used the AMR capabilities of the code, which guarantee better resolution where needed and a better usage of computational resources.

Our simulation has been run with a cubic domain of $863 R_{\odot} = 4 \text{ AU}$ on a side and a coarse grid resolution of 128 cells per side. We adopt two levels of refinement with a refinement factor of two (i.e., when a cell is refined it is divided by two along each dimension), in this way the smaller cell size is $1.68 R_{\odot}$, as was the case in the 256³ simulations of P12. The refinement criterion is based on cell gas density. Cell densities above $1.38 \times 10^8 \text{ g cm}^{-3}$ dictate a cell division. Additionally ENZO adaptively de-refines the zones where a cell and its surrounding region no longer satisfy the refinement criterion. For our choice of the smoothing length (see below), two levels of refinement are the minimum to obtain a stable giant model with the best possible energy conservation. ENZO has various numerical solvers for hydrodynamics, for all our simulations we adopt the ZEUS solver, an implementation of the one used by the code ZEUS (Stone & Norman 1992a, Stone & Norman 1992b). The value of the artificial viscosity we use (roughly representing the number of cells over which a shock is spread) is 2.

As we will explain in Section 3.2.1 and Section 3.2.2, we use point-masses, interacting only gravitationally with both gas and other particles, to model the primary core and the companion. These point-masses have a smoothing length associated with them, according to the prescription of Ruffert (1993), and in this work with the term “smoothing length” we refer to the smoothing length only of the point-mass particles. To ensure a reasonable degree of energy conservation in the simulations, we use a smoothing length equal to 3 times the smallest cell size. This was found

to be the optimal value by Staff et al. (2016a), who monitored the energy conservation in their CE ENZO simulations as a function of smoothing length¹. The smoothing length is described in detail in Chapter 6.

The methodology followed to simulate our CE interaction consists of two main phases and is described in the following sections.

3.2.1 Single star setup and stabilisation

As in P12 we model our binary system as an RGB primary and a smaller companion with comparable mass, identifiable with a main sequence star or a compact object such as a white dwarf. The resolution is not sufficient to resolve the primary’s core, nor the companion, so we model them as dimensionless point-masses. The companion mass is $M_2 = 0.6 M_\odot$ (this choice will be discussed in Section 3.2.2). The primary star is an extended object whose envelope is well resolved. We use the same initial model as in P12: a star with an initial mass of $1 M_\odot$ evolved to the RGB with the 1D stellar evolution code EVOL (Herwig 2000). At this stage of the evolution the star has a radius of $R_1 = 83 R_\odot$, a total mass of $M_1 = 0.88 M_\odot$ and a core mass of $M_c = 0.392 M_\odot$.

The relevant ENZO physical quantities are interpolated from the 1D model to the 3D domain. The star is stabilised in the domain, as described in Section 2.2, where we carried out the damping for $10 t_{\text{dyn}}$ time and the stabilisation for an additional 10. At the end of this process the initial 3D stellar model is relaxed with respect to the 1D model as showed in Figure 3.1. The sharp density jump at the edge of the star has been smoothed by the stabilisation process, and the star is now slightly larger. The contour of density of $10^{-11} \text{ g cm}^{-3}$ has a radius $100 R_\odot$. The central density is also slightly reduced, but overall the original structure of the star is mostly preserved.

As an additional test of the stability (see also Section 2.2), the velocities that develop have been compared to global and local velocity scales, such as the local sound speed and the dynamical velocity, $v_{\text{dyn},1} = R_1/t_{\text{dyn},1} \simeq R_1(G\langle\rho_1\rangle)^{\frac{1}{2}}$, where $t_{\text{dyn},1}$ is the dynamical time of the primary, G is the gravitational constant and $\langle\rho_1\rangle$ is the average density of the star. Additionally, we also compare the gas velocities in the frame of reference of the primary to the orbital velocities of the binary

¹I am co-author of the paper. The work I have done for it consisted of quantifying how non-negligible approximations in the numerical solver of Enzo result in reduced conservation due to poorly resolved steep, moving density gradients.




FIGURE 3.1: Radial density profiles of the primary RGB star used in our simulation, calculated with the 1D EVOL code (blue curve), after mapping it in the ENZO computational domain but before the stabilisation process (red curve) and after stabilisation (yellow curve). The change in slope at a radius of $3 \times 10^{-2} R_{\odot}$ marks the core-envelope boundary of the 1D model, while the vertical line shows the size of an ENZO cell at the deepest level of refinement.

system in the frame of reference of the center of mass (see Section 2.2). The star was relaxed for 10 dynamical times and mostly maintained its internal structure (Figure 3.1). At each step during the relaxation at most 7 per cent of the cells had velocities exceeding the lowest of the velocity limits discussed above. Hence we expect the contamination of the CE interaction by the spurious motions of the primary envelope to be negligible.

3.2.2 Binary system setup

The companion has a mass $M_2 = 0.6 M_{\odot}$, selected among those simulated by P12, also based on the fact that their $0.6 M_{\odot}$ companion simulations were converged for the coarse grid resolutions we are using. The orbital separation was the largest that would result in the evolution of the orbital elements and eventually in a CE within a reasonable computational time: $a = 300 R_{\odot}$ (corresponding to a period of 496 days = 1.36 yr). This value also corresponds to the approximate maximum orbital separation from which a tidal capture of the companion may take place within the evolution of a star similar to our primary (Madappatt et al. 2016). The system was placed in circular orbit, where we gave the RGB star a Keplerian velocity $v_1 \simeq 12.4 \text{ km s}^{-1}$ and the companion point particle a velocity $v_2 \simeq 18.2 \text{ km s}^{-1}$, with the point mass core of the primary coinciding with the centre of the box. The initial configuration of the binary system is shown in




FIGURE 3.2: Density slice on the $x - y$ plane at $z = 0$, showing the entire computational domain with the initial setup of the binary simulation. The arrows show the velocity field; vectors are scaled to the maximum value of the velocity (the orbital velocity of the companion or 18.2 km s^{-1}). Note that also the point-mass particles representing respectively the core of the primary and the companion are given Keplerian velocities, even if arrows are not present in the plot.

Figure 3.2.

In our simulation the primary is driven into Roche lobe contact (the Roche lobe radius of the primary is $124 R_{\odot}$ at an orbital separation of $300 R_{\odot}$, using the approximation of Eggleton 1983, but noting it to be valid in the case of synchronised orbits, which is not our case) and eventually a CE interaction by the pre-contact tidal interactions in a relatively short time-scale, much shorter than realistic tidal interaction time-scales. The reason for this difference is that the strength of the interaction is sensitive to departures of the stellar envelope distribution from spherical symmetry. Inserting the companion in the computational domain generates a small distortion of the primary's envelope resulting in a set of oscillations, which exert a relatively strong tidal force. Paradoxically, this larger than average tide results in shortening of the orbital separation within reasonable computational times, something that would not be so if the tide were better reproduced.

We do not apply any initial rotation to the primary. However, we achieve a spinning star by spin-orbit interaction. The initial orbital separation is the maximum that would allow a capture of the companion into a CE interaction during this phase of the evolution of the star. This means that




FIGURE 3.3: Left panel: evolution of the separation, a , between the two particles representing the core of the primary and the companion, over the whole simulation time. The blue line represents the actual separation computed every 0.01 year. The red line represents the separation averaged over one orbital cycle. The black vertical lines represent, from left to right, the beginning of mass transfer, the beginning of the fast in-spiral phase and the end of the fast in-spiral phase. Right panel: evolution of the orbital decay, computed on the separation averaged over one orbital cycle.

the total angular momentum in the system, which is increasingly transferred from the orbit to the envelope of the giant, is approximately that which would be expected for this system.

3.3 Results

3.3.1 Orbital Separation

The separation between the point masses as a function of time and the orbital decay rate are shown in Figure 3.3. To determine the time when the mass transfer phase begins, we calculated the Roche lobe surface around the primary using the total potential field computed in the simulation. Then, we checked whether the cells contained within the primary's Roche lobe, including the first cell near the inner Lagrangian point in the companion's Roche lobe, have a density greater than the vacuum's density ($6.93 \times 10^{-12} \text{ g cm}^{-3}$). Computed in this way, the beginning of the contact phase takes place after about 547 days $\approx 1.5 \text{ yr}$ from the beginning of the simulation. During this pre-contact phase, the orbital separation has been reduced from 300 to 265 R_{\odot} , at which point the primary's Roche lobe radius is 108 R_{\odot} , similar to the stellar radius at the start of the simulation.

The mass transfer phase lasts until the companion is engulfed in the envelope of the primary,

at which point the rapid in-spiral phase begins. We define the start of the rapid in-spiral phase as the time when the equipotential surface passing through the outer Lagrangian point L_2 has a density greater than the vacuum's density in each of its cells. This condition is satisfied after about 1515 days or 4.2 yr from the beginning of the simulation.

The rapid in-spiral phase is observed as a steepening of the separation vs time curve, which denotes a regime change. This phase lasts 324 days and ends at 1840 days, or 5.0 yr from the beginning of the simulation, when the orbital separation stabilises. We have used the same criterion as P12 and Sandquist et al. (1998), who defined the end of the rapid in-spiral phase when $-\dot{a} < 0.1(-\dot{a}_{\max})$, where $\dot{a} = da/dt$. This point is somewhat arbitrary because it depends on how steep the in-spiral is. In our simulations, the in-spiral is much steeper than that witnessed in the simulations of Sandquist et al. (1998) and P12, as can be seen by comparing our Figure 3.3, lower panel with their figures 4 and 5, respectively.

The rapid in-spiral phase in our simulation lasts approximately 10 per cent longer than for the equivalent simulation of P12, and longer still if we acknowledge that at the end of the in-spiral phase as defined above, the separation is still reducing considerably. This could be due to the fact that our donor star is puffed up by the interactions in the previous phases, hence it is less dense. The delayed rapid in-spiral and its longer duration are in line with the results obtained by P12 in their simulations with the companion star slightly away from the primary surface rather than in contact. The rapid in-spiral time-scale is 65 per cent of the initial period, in line with theoretical expectations (Paczynski 1976).

The orbit starts to become elliptical during the rapid in-spiral phase. Using the maxima and minima in the orbital separation evolution after the end of the rapid in-spiral phase, we obtain an eccentricity $e = 0.12$, in agreement with what was obtained by P12.

The final separation achieved (a_f) is a crucial output of the CE simulations. P12 identified that CE simulations have final separations that not only tend to be larger than observed (Zorotovic et al. 2010 and De Marco et al. 2011), but that depend on the companion/primary mass ratio (q), a tendency not seen in the observations. By using the average separation (red line in Figure 3.3) we estimated the value of the separation reached at the end of the rapid in-spiral phase to be $36 R_\odot$, using the criterion described above, and $20 R_\odot$ if we take the average value at the end of the simulation (see Table 3.1, where we report the initial conditions and final outcomes for all past CE




FIGURE 3.4: Final orbital separation vs mass ratio $q = M_2/M_1$ for observed post-CE systems (Zorotovic et al. 2010 and De Marco et al. 2011, black dots) and for simulations (Sandquist et al. 1998, green circles; Ricker & Taam 2012, cyan triangle - note that here we report the separation of the simulations of Ricker & Taam 2012 which is lower than reposted in Ricker & Taam 2008 where the in-spiral had not come to an end; the 256^3 ENZO simulations of P12 are shown as yellow squares; Rasio & Livio 1996, magenta pentagon; PHANTOM simulations, blue hexagons, see Section 3.4; Ohlmann et al. 2016, grey diamond; Nandez et al. 2015, pink cross), together with the results of our new ENZO simulation (yellow star). Note that for our work and for the simulations of Sandquist et al. (1998) the definition of the final separation may lead to a slightly larger value than for other simulations. See text and Table 3.1.

simulations including at least a giant). The separation at the end of the simulation is $\simeq 4$ times the smoothing length, indicating that the end of the in-spiral is not affected by the smoothing-length and resolution. Our values of the final separation are larger than those of P12, which were 19 and $16 R_\odot$, for the criterion-defined and at the end of the simulation, respectively. In other words, the final separation is larger by 25 per cent for the ENZO simulation starting with a larger initial separation. We plot final separations in Fig. 3.4, and we defer further discussion on the issue of final orbital separation to Section 3.5.

M_1 (M_\odot)	$M_{1,c}$ (M_\odot)	R_1 (R_\odot)	Giant	M_2 (M_\odot)	q	a_i/R_1	Ω/ω^1	$R_1/R_{1,RL}$	Code ²	Resolution (Part./ R_\odot)	τ_{run}^3 (day)	a_f^4 (R_\odot)	M_{Unb}^5 (%)	Ref. ⁶
4	0.7	66	RGB	0.7	0.18	1.6	1(y)	1.3	SPH	500k	124	1(e)	10(?)	1
3	0.7	200	AGB	0.4	0.13	1.4	1(n)	1.3	n-grid	2.4	800	4.4(a)	41(?)	2
3	0.7	200	AGB	0.4	0.13	1.4	0	1.3	n-grid	2.4	800	4.7(a)	46(?)	2
5	1.0	200	AGB	0.4	0.08	1.4	1(n)	1.2	n-grid	2.4	800	4.4(a)	21(?)	2
5	1.0	200	AGB	0.6	0.12	1.4	1(n)	1.3	n-grid	2.4	800	4.8(a)	45(?)	2
5	0.94	354	AGB	0.6	0.12	1.5	0	1.2	n-grid	2.4	800	8.9(a)	46(?)	2
1.05	0.36	31	RGB	0.6	0.57	2.0	0.95(n)	1.2	a-grid(F)	0.29	60	9(e)	26(t)	3
0.88	0.39	85	RGB	0.1	0.11	1.0	0	1.8	u-grid(E)	1.7	1000	5.7(a) / 4.2(e)	–	4
0.88	0.39	85	RGB	0.15	0.17	1.0	0	1.9	u-grid	(E)1.7	1000	6.9(a) / 4.7(e)	–	4
0.88	0.39	85	RGB	0.3	0.34	1.0	0	2.1	u-grid(E)	1.7	1000	11(a) / 9.0(e)	–	4
0.88	0.39	85	RGB	0.6	0.68	1.0	0	2.4	u-grid(E)	1.7	1000	19(a) / 16(e)	–	4
0.88	0.39	85	RGB	0.9	1.02	1.0	0	2.6	u-grid(E)	1.7	1000	26(a) / 22(e)	–	4
0.88	0.39	83	RGB	0.1	0.11	1.0	0	1.8	SPH(S)	500k	1050	6.1(a) / 5.7(e)	2(t)	4
0.88	0.39	83	RGB	0.15	0.17	1.0	0	1.9	SPH(S)	500k	950	7.3(a) / 7.8(e)	6(t)	4
0.88	0.39	83	RGB	0.3	0.34	1.0	0	2.1	SPH(S)	500k	750	11(a) / 10(e)	8(t)	4
0.88	0.39	83	RGB	0.6	0.68	1.0	0	2.4	SPH(S)	500k	950	21(a) / 18(e)	10(t)	4
0.88	0.39	83	RGB	0.9	1.02	1.0	0	2.6	SPH(S)	500k	600	27(a) / 25(e)	10(t)	4
1.98	0.38	49	RGB	0.99	0.5	1.0	0.95(n)	2.3	m-mesh	0.07-0.01	120	4.9(e)	8(t)	5
1.50	0.32	26 ⁷	RGB	0.36	0.24	2.0	0	1.0	SPH	200k	(?)	0.91(e)	100(r) ⁸	6
0.88	0.39	100.	RGB	0.6	0.68	3	0	0.81	a-grid(E)	1.7	2000	36(a)/20(e)	16(t)	7
0.88	0.39	93	RGB	0.6	0.68	1.1	0	2.2	SPH(P)	2.3m	1850	21(a)/16(e)	12(t)	7
0.88	0.39	91	RGB	0.6	0.68	2.4	0	1.0	SPH(P)	300k	5050	29(a)/22(e)	15(t)	7

¹Stellar spin frequency as a function of orbital frequency, with an indication of whether the star was stabilised in its rotating configuration (y) or not (n) before the start of the simulation.

²SPH: smooth particle hydrodynamics; u-grid: static uniform grid; n-grid: static nested grids; m-mesh: moving mesh; a-grid: adaptive mesh refinement grid; F: FLASH, E: ENZO, S: SNSPH, P: PHANTOM.

³ Information not provided (?).

⁴Rounded to 2 significant figures, calculated either at the end of the simulation (e) or at a time defined by the formula in Section 3.3.1 (a).

⁵ Calculated by including thermal energy (t), not including thermal energy (k), information not provided (?) or including recombination energy (r).

⁶1: Rasio & Livio 1996. 2: Sandquist et al. 1998. 3: Ricker & Taam 2012. 4: P12: Passy et al. 2012. 5: Ohlmann et al. 2016.

6: Nandez et al. 2015. 7: This work.

⁷ This is the Roche lobe radius also corresponding to the SPH radius in their simulation.

⁸ Note that the same simulation run without recombination energy unbinds 50 per cent of the envelope, although the authors of that simulation do not present data to illustrate their statement.

TABLE 3.1: A comparison of initial conditions and final outcomes of previous common envelope simulations that included at least one giant star.

3.3.2 Envelope ejection

To determine the extent to which the envelope is unbound we determined whether gas has total energy larger than zero. The total energy can be calculated including or excluding thermal energy, where the former prescription results in more unbound gas. Ivanova & Chaichenets (2011) discussed how it is the enthalpy rather than the thermal energy that needs to be included when determining whether a gas parcel is bound or not. Using enthalpy instead of thermal energy increases the unbound mass very marginally.

In Figure 3.5 we present density slices in the orbital and perpendicular planes. In the first and middle columns we compare the distribution of unbound gas both including thermal energy (left column) and excluding thermal energy (middle column), to distinguish between gas acceleration and gas heating. The initial unbinding event (first two rows, left columns) happens because of heating of the gas falling into the potential well of the companion during the mass transfer phase, which is why this unbound material is not recorded on Figure 3.5, middle column. This unbound material has very low mass. The heating of the gas in the unbound zones can be seen in Figure 3.6, where we plot the temperature distribution. Note also that the very hot zone outside of the envelope is due to the presence of the low-pressure, high-temperature vacuum. Later, during the rapid in-spiral phase (Figure 3.5, last two rows, left and middle columns) far more mass is unbound because it is accelerated above the escape velocity as demonstrated by the similarity of the left and central columns. In this case this behaviour can be seen by plotting the velocity distribution (Figure 3.7). The masses of the unbound gas inside the simulation box obtained by including or excluding E_{therm} are listed in Table 3.2.




FIGURE 3.5: Left panel, left column: density slices perpendicular to the z axis in the orbital plane after (from top to bottom) 887, 1381, 1669 and 1724 days from the beginning of the simulation. The point-mass particles representing the core of the primary and the companion are shown as black dots, while the white regions represent the unbound gas. The size of the black dots is not representative of any property of the point masses and is chosen only to highlight them. Left panel, right column: same as the left column, but excluding thermal energy (E_{th}) in the computation of the bound/unbound mass elements. Right panel, left and right columns: same as for the left panels, but the density slices are taken perpendicularly to the orbital plane, at $x = 0$.




FIGURE 3.6: Same as the left and right panels of Figure 3.5, but the showing the temperature distribution instead.




FIGURE 3.7: Same as the left and right panels of Figure 3.5, but the showing the velocity distribution instead.

The evolution of the unbound gas can be followed only inside the simulation box, due to the grid nature of ENZO. However, we estimated whether the mass that leaves the box is bound or unbound in the following way. We calculated the fraction of unbound gas contained within the box boundary (i.e., within the six, one cell thick, box faces) and we assumed it to be representative of the fraction of unbound gas between code outputs (which take place every 3.65 days = 0.01 yr). We then multiplied this fraction by the mass that leaves the box between code outputs.

The procedure we adopted above is computationally cheap, since it is entirely done during post-processing of the simulation data. However, to estimate the goodness of this procedure, we performed a one-off simulation using the “inline” *Python* capability of ENZO. This option allows us to measure user-defined quantities at each coarse grid time-step of the simulation at the expense of simulation time and works by coupling ENZO with the YT data analysis software (Turk et al. 2011). With this option turned on, ENZO stores its relevant arrays in memory in such a fashion that YT can read them as they were saved on the disk, then executes a script (written by the user) that analyses the data. More detailed information on how this works can be found in Chapter 2 and Bryan et al. 2014.

In our script we select all cells defining the six faces of our cubic simulation box. For each of these faces we then compute the average component of the velocity directed perpendicularly outward, $\langle v_{\text{out}} \rangle$, for all those cells having a velocity component directed outwards (whose fraction is f_{out}). With these quantities it is possible to obtain an estimate of the fractional volume of the face that will be leaving the box during the next time-step as

$$V_{\text{lost}} \simeq s \langle v_{\text{out}} \rangle f_{\text{out}} \Delta t , \quad (3.1)$$

where s is the surface area of the face and Δt is the time-step.

Finally, from this value one can estimate the bound and unbound mass lost from the box as

$$\begin{cases} m_{\text{lost,bound}} \simeq \frac{V_{\text{lost}}}{V} m f_{\text{bound}} \\ m_{\text{lost,unbound}} \simeq \frac{V_{\text{lost}}}{V} m (1 - f_{\text{bound}}) \end{cases} , \quad (3.2)$$

where m is the total mass in the face, V is the face volume and f_{bound} is the fraction of bound cells,

determined by considering the cell’s total energy according to one of the prescriptions discussed in Section 3.3.2. This procedure has the advantage of recording the mass losses from the domain at each coarse-grid time-step, therefore with a cadence abundantly superior with respect to the one with which data are saved to disk.

It is important to remark that the simulations are already slowed down when the “inline” *Python* option is activated. In addition the particular script we run is especially expensive in terms of computational time, this makes running every simulation with this option turned on beyond our computational possibilities. For this reason we only run one simulation of this type, with the same setup of the one described in the published paper, but with uniform, static grid to further reduce the load and the time used. We then applied the approximated method on the post-processed data and compared the results. The outcome of the comparison showed that the bound and unbound fractions are very similar using either the “inline” *Python* or interpolating between the dumps saved on disk by the code. Therefore we used the second procedure on the published data.

In Figure 3.8 we present the various components of the mass (total, bound and unbound) for both the gas inside the simulation domain (left panel) and the gas that accumulates outside of the domain (right panel). Looking at the latter it is possible to see that our approximation is consistent with the total amount of mass that leaves the box during the simulation, shown in Figure 3.8 (left panel). The first unbound mass leaves the box at approximately 1500 days, at the onset of the rapid in-spiral, but the bulk of the mass flows out during the rapid in-spiral phase (between approximately 1750 and 1900 days). The total mass unbound in the simulation amounts to $8 \times 10^{-2} M_{\odot}$, or 16 per cent of the initial envelope mass. The unbound mass is 14 per cent, if we do not include thermal energy and 17 per cent, if we use the enthalpy as suggested by Ivanova & Chaichenets (2011). P12 found that 10 per cent of the initial envelope mass was unbound, which should be compared to our 16 per cent. This increase likely represents the effect of a larger initial separation.

Most of the ejecta is expected to flow away close to the orbital plane, where the gas is accelerated by the orbiting particles. This was already borne out by the simulations of Sandquist et al. (1998) and is clearly seen in Figure 3.5. Figure 3.9 demonstrates how the envelope is ejected around the binary over time. We divide the computational domain into six pyramids centred at the centre of the box and whose bases are the six faces. We plot the mass in pairs of pyramids aligned with each of the three directions, x , y and z . Initially the mass is equally distributed in the three




FIGURE 3.8: Left panel: evolution of the gas mass inside the simulation box over time. The blue line represents the value of the initial gas mass contained in the box and is plotted for comparison, while the green line shows the evolution of the total mass contained inside the box. The red and cyan lines show, respectively, the bound and unbound components of the mass. Right panel: cumulative mass of the gas flowing out of the simulation box over time. Line colours have the same meaning as for the upper panel. The black vertical lines in both panels correspond to the beginning and end of the rapid in-spiral and both the plots are limited to the part of the simulation where significant mass is lost from the box.

pairs of pyramids as the star resides at the centre of the box. Later the mass in the pyramid pairs oscillates as the giant moves along its orbit. The decrease of the peaks in the green line during the fast in-spiral phase in Figure 3.9 marking approximately the completion of a full orbital revolution, demonstrating a decrease in the mass contained in the z direction in favour of mass contained in the other two directions. The decreasing amplitude of the oscillations over time indicates that the gas distribution becomes more and more independent of the orbital motion of the two particles, as

Time (days)	$M_{\text{unb}} (E_{\text{kin}} + E_{\text{pot}} + E_{\text{therm}})$ (M_\odot)	$M_{\text{unb}} (E_{\text{kin}} + E_{\text{pot}})$ (M_\odot)
887	1.61×10^{-3}	7.64×10^{-6}
1381	1.58×10^{-3}	4.89×10^{-5}
1669	1.56×10^{-2}	6.28×10^{-3}
1724	3.72×10^{-2}	2.93×10^{-2}

TABLE 3.2: Gas mass unbound inside the simulation box at various times and for the two possible prescriptions, including thermal energy and excluding thermal energy.




FIGURE 3.9: Gas mass inside the simulation domain vs time for the gas located in six pyramids whose bases are the six faces and whose vertexes are at the centre of the domain. The two pyramids along the x axis are in blue, along the y axis are in red and along the z axis are in green. The cyan line shows the sum of the x and y contributions to highlight the behaviour of the mass ejection in the orbital plane. The black vertical lines show the estimated beginning of the Roche-lobe overflow phase and the beginning and end of the fast in-spiral phase.

the interaction proceeds. Towards the end of the CE, as the oscillations cease, more mass is being ejected out of the simulation box highlighting how the fast in-spiral rapidly lifts the envelope, disrupting the primary star.

3.3.3 Tidal bulges

As explained in Section 3.2.2, the pre-contact phase in our simulation takes place over much shorter time-scales than it would in nature. The short pre-contact time-scale observed in our simulation is due to deformations created on the primary by the insertion of the companion into the computational domain that are likely larger compared to those expected from the theory of equilibrium tides.

A simple analytical estimate of the mass, δM_1 , contained in the tidal bulges of the primary, for equilibrium tides, can be obtained from Zahn (2008):

$$\delta M_1 \leq M_2 \left(\frac{R_1}{a} \right)^3, \quad (3.3)$$

where M_1 , M_2 , R_1 and a are the masses of the primary, secondary, the radius of the primary and the orbital separation, respectively. For the purpose of this calculation we only vary a with time,




FIGURE 3.10: Mass in the tidal bulges of the primary star overtime during the pre-contact phase, estimated from the ENZO simulation data (solid blue line), from the PHANTOM simulation data described in Section 3.4 (dashed blue line) and from the analytical formula (dotted red line).

while leaving R_1 constant. We compared this analytical estimate with the bound mass residing outside the initial equilibrium radius of the primary.

As shown in Figure 3.10 (solid blue line), the insertion of the companion into our ENZO simulation triggers some oscillations, on a time-scale of the order of the dynamical time of the star ($\simeq 21$ days). Over the pre-contact phase there is also a gradual expansion of the star, seen as an increasing trend of the mass outside its original volume. Once mass transfer starts at $\simeq 550$ days, new forces are at play that overwhelm the tidal torque. Figure 3.10 also shows that the mass outside the original stellar volume in the ENZO simulations does not exceed substantially the theoretical value (note that the dashed dotted line is calculated using the PHANTOM simulations presented in Section 3.4 and will be discussed at that time). Its distribution is in the shape of two small opposite bulges that are initially aligned with the direction of the companion, which disappear and reappear at 90 degrees to the original direction. This generates the relatively strong torques that result in the fast decrease of the orbital separation during the pre-contact phase.

3.3.4 Evolution of the gas velocities and density in proximity to the companion: the end of the in-spiral

The mechanism behind the energy and angular momentum exchange that drives the in-spiral is gravitational drag (Ricker & Taam 2012). Gravitational drag is caused by the gas which flows past the moving body (in our case the companion star), forming a wake with higher density behind it that gravitationally pulls on it, slowing the body down. The gravitational drag experienced by a body immersed in a fluid depends on the body’s mass, the fluid density, the velocity contrast between the body and the fluid and on the Mach number of the body. Approximations for the gravitational drag are given by Iben & Livio (1993, $F_{drag} \propto (M_2 \rho v_{rel}^2) / (v_{rel}^2 + c_s^2)$, for the subsonic motion regime) and by Ostriker (1999) who calculated a more detailed formula, carefully considering the effects of the Mach number.

It is fundamental to determine whether simulations accurately reproduce the effects of gravitational drag because this determines in turn when the companion in-spiral terminates and, as a result, the amount of orbital energy deposited. Is the end of the in-spiral due to the decreasing density around the particles, the co-rotation of the surrounding gas or a change in the Mach regime (as was the case in the simulations of Staff et al. 2016b)? Does the density gradient affect the force as questioned by MacLeod & Ramirez-Ruiz (2015)? How does the interplay of resolution and smoothing length affect the simulation (Staff et al. 2016a)? It is well known that the particles will not approach closer than approximately two smoothing lengths, effectively because their potentials are flat within that distance. However, less clear are the effects that not resolving a radius of the order of the Bondi radius (Bondi 1952) around the particles will have on the drag force (Staff et al. 2016b). Finally, the gas trapped in the potential of the particles and travelling with them, gives the particles a larger mass and an effective “size” and this could affect both the gravitational and hydrodynamic drags.

In Figure 3.11 we display the evolution of the density profile between the two cores, showing only the part between the particles (left panel), or the entire computational domain (right panel). The density profile changes smoothly at the beginning of the simulation, with the primary expanding, but it then transitions into a phase of more rapid change at the onset of the rapid in-spiral phase, when the profile flattens and then becomes U-shaped, showing peaks at the locations of




FIGURE 3.11: Left panel: density profile between the core of the primary (located at zero in the abscissa) and the companion (each black dot represents the density at the location of the companion). The colour legend shows the time, in days, corresponding to each line. Profiles are taken, for clarity, every 110 dumps of the code. Right panel: same as the top panel, but extended to the whole box. The primary’s core is represented by a large black dot while the companion is marked as a smaller dot.

both the primary’s core and the companion with densities of $2.8 \times 10^{-6} \text{ g cm}^{-3}$ for the primary and $4.6 \times 10^{-6} \text{ g cm}^{-3}$ for the companion. The underlying density is of the order of $10^{-6} \text{ g cm}^{-3}$ at 365 – 730 days after the start of the simulation. These values are comparable to those of P12 (their figure 13, middle panel).

The gas density in the proximity of the particles at the end of the simulation is high, and is unlikely to be the cause of the observed slowing down of the in-spiral. From the density profiles Figure 3.11 (right panel) it is clear that during the evolution of the system some of the envelope accumulates around the companion. The accumulation of mass is negligible until the beginning of the rapid in-spiral phase, during which it starts to increase because the companion is plunging into the denser parts of the envelope. The companion local density is a factor of a few larger than the density 10-20 R_{\odot} away from it. The density gradient underlying the density peak near the companion is small and likely unimportant to the in-spiral.

In Figure 3.12 (left panel) we plot the companion’s speed, the average local gas velocity projected in the direction of motion of the companion and the average local gas velocity projected in the direction perpendicular to the motion of the companion, while in the right panel we plot the companion’s Mach number and the normalised average density near the companion. To calculate the parallel and perpendicular ambient gas velocities we averaged the respective projections for all

cells within a volume with radius $10 R_{\odot}$ from the companion. The local density was calculated by averaging the density inside the same volume and the Mach number by averaging the gas sound speed within the same volume.

As was the case for the simulation of P12, the entire journey of the companion is subsonic, reaching at most a Mach number of 0.6. This is different from the simulations of Staff et al. (2016b), where the initial part of the in-spiral was supersonic and the end of the in-spiral phase appeared to coincide with the transition between a super-sonic and sub-sonic regimes. No such transition occurs here. In our simulation no shocks are observed and no shock heating can take place as was instead the case in the simulations of Sandquist et al. (1998).

On the other hand, a regime change does take place at the approximate time of the end of the in-spiral, at ~ 1850 days. The density peaks at that time, due to the fact that first the companion plunges into a region of higher density, but as the in-spiral stalls, that volume is somewhat evacuated as gas flows outwards. The velocity of the companion, which increases during the in-spiral due to the decreasing orbital separation, approximately levels off after the in-spiral. The peaks and troughs are due to the orbit becoming eccentric early in the in-spiral. The local gas velocity parallel to the companion's motion increases as the local gas is dragged along with the companion. Corresponding with the approximate end of the in-spiral the local gas is spun up to almost co-rotate with the companion, which greatly reduces the drag. At that time the gas perpendicular velocity is seen to increase somewhat, as an outflow takes place (which acts to decrease the local density). During the in-spiral, the perpendicular gas velocity local to the companion decreases overall, as the companion plunges into layers that are harder and harder to eject. The peak at 1700 days corresponds with the high companion velocity peak taking place at the first periastron of the now-eccentric orbit. At the same time, much of the unbinding is taking place at the location of the companion, as can be seen in Figure 3.5 (third row).

To confirm that this trend is not a result of the size of the sphere used to estimate our quantities, we carried out the same test with spheres of $5 R_{\odot}$ and $20 R_{\odot}$. Both show results similar to Figure 3.12 with the only exception that the gas velocity parallel to the companion direction of motion is overall larger and close to the companion's velocity for the smaller sphere, as expected. We also note that at the beginning of the in-spiral the local gas has a rotation velocity of $10\text{-}20 \text{ km s}^{-1}$, which is a range of values expected for giants spun up by a companion.




FIGURE 3.12: Left panel: companion velocity (thicker blue line), local average gas velocity projected on the direction of the companion velocity ($\langle v_{gas,\parallel} \rangle$, thick green line) and local average gas velocity perpendicular to the direction of the companion velocity ($\langle v_{gas,\perp} \rangle$, thin cyan line). The three lines are smoothed with a Savitzky-Golay filter, using 31 coefficients and 7th order polynomials. Right panel: companion Mach number (thick blue line) and normalised average gas density in the companion's proximity ($\langle \rho \rangle / \langle \rho \rangle_{max}$, thin green line). All plots start at the onset of the rapid in-spiral, the vertical solid lines represent the estimated end of the rapid in-spiral and the dashed ones mark the point of maximum density.

3.3.5 Angular momentum and energy conservation

Energy and angular momentum were excellently conserved by the SPH simulations of P12 (to a 1 percent level). They did not check the conservation level of their equivalent ENZO simulations, because of the grid nature of the code which leads to loss of mass off the simulation box and because their ENZO simulations showed similar results to the SPH ones, which implied a reasonable level of energy conservation.

As introduced in Section 3.2, Staff et al. (2016) quantified the level of energy non-conservation in grid based simulations using ENZO and determined that conservation is improved by selecting a larger smoothing length of 3 cells rather than what was used by P12 (1.5 cells, see also Chapter 6). The highest resolution in our AMR simulation is the same as the resolution in the unigrid simulations of P12. However, we have adopted the larger smoothing length of 3 cells, which must have weakened the gravitational interaction somewhat compared to the simulations of P12.

In Figure 3.13 and 3.14 we plot various components of the angular momentum and energy, respectively, in the computational domain as a function of time. The behaviour of some of the components is driven by mass loss out of the computational domain, which starts at ~ 260 days




FIGURE 3.13: Evolution of the angular momentum components with respect to the center of mass of the system for gas, particles and their sum, inside the simulation domain. The blue curves represent the x angular momenta, the red curves the y angular momenta and the green curves the z angular momenta. In all three cases the solid line represents the particles + gas, the dashed lines the particles and the dotted line the gas. The black vertical line represents the moment when the envelope mass starts leaving the box ($\simeq 260$ days).

(some of the low density ambient medium outflows before, but has negligible mass), but is particularly heavy during the rapid in-spiral phase. In Figure 3.13 we see that, as expected, the z component of the angular momentum dominates over the other components and that most of the angular momentum resides in the point masses, with an initial value of $\sim 3.5 \cdot 10^{52} \text{ g cm}^2 \text{ s}^{-1}$. Before 260 days from the beginning of the simulation, only negligible mass and angular momentum are leaving the simulation box. The particles' z angular momentum decreases during the in-spiral. Some of that is transferred to the gas. Five percent of the angular momentum is lost due to non-conservation, between the beginning of the simulation and 260 days, while 10 percent is lost over the first 3 years, a time at which substantial amount of mass starts leaving the box. This value is larger (as expected) than for the SPH simulation of P12 and similar to the 8 per cent of Sandquist et al. (1998), who estimated it over $\simeq 1000$ days of their simulation.

Estimating the level of conservation of energy is even more difficult than for the angular momentum, because the low density medium filling the volume outside the star has a very high thermal energy, even if its total mass is negligible. Even before envelope mass starts flowing out of the computational domain at 260 days, a small amount of this high energy gas flows out of the box




FIGURE 3.14: Components of the energy as a function of simulation time in the domain: total energy (thick black line), total kinetic energy (solid black line), total potential energy (dashed black line), total (= gas) thermal energy (dotted black line), gas kinetic energy (solid red line), gas potential energy (dashed red line), point-mass kinetic (solid yellow line), point-mass to point-mass potential (dashed yellow line) and point-mass to gas potential (dashed cyan line). The black vertical line represents the moment when the envelope mass starts leaving the box ($\simeq 260$ days).

taking with it an energy of $\simeq 1.3 \times 10^{45}$ erg (or $\simeq 11$ percent of the initial total energy). This behaviour is clear in Figure 3.14: the total energy at the beginning of the simulation is dominated by the thermal energy of the “vacuum” and by the potential energy between the point mass particles and the gas, with the former continuously decreasing as some of the low density medium flows out of the box; this decrease is mimicked by the total energy at times greater than 260 days. Before this threshold is passed the code conserves energy to the 4 percent level, similar to the result of Sandquist et al. (1998).

3.4 Comparison with PHANTOM simulations

Alongside the ENZO simulation we carried out a set of simulations using the SPH code PHANTOM (Price & Federrath 2010; Lodato & Price 2010; Chapter 2).

In PHANTOM we map the same 1D stellar model used for the ENZO simulations, but in this case the SPH particles are distributed so as to reproduce the entire stellar mass distribution, inclusive of the core. This generates a very high particle density at the location of the core that would slow down the simulation excessively. Therefore, similarly to the ENZO procedure, we use point-masses (called sink particles in the phantom nomenclature), interacting only gravitationally with

both gas and SPH particles, to model the primary core and the companion. These sink particles were made to accrete all SPH particles within a radius of $0.03 R_{\odot}$. This quickly generated a “core” with a mass of $M_c=0.392 M_{\odot}$. The giant was stabilised as was done for the ENZO simulations. The companion particle was placed in the computation with a mass of $0.6 M_{\odot}$. Both core and companion particles were given a softening length² of $3 R_{\odot}$. In all our PHANTOM simulations we use a default number of neighbours equal to $\simeq 60$, adaptive softening for the SPH particles and a Courant factor of 0.4. For the shock capturing we use an artificial viscosity equal to 10% of the particles’ sound speeds. Finally, note that the number of particles mentioned for all the phantom simulations in the following sections is the actual number of particles after the accretion process (e.g., the convergence test using 2.3×10^6 particles (see below) was actually initialised with 4×10^6 particles). The giant was then damped and stabilised as was done for ENZO.

Our first simulation, using 2.3 million particles, has similar parameters to that carried out by P12 with a companion mass of $0.6 M_{\odot}$ and is used as a verification step to ensure that PHANTOM performs similarly to ENZO and the SPH code used by P12, SNSPH (Fryer et al. 2006). This simulation’s outcomes were compared directly with the SPH simulation “SPH2”, which in that study was carried out with the SPH code SNSPH using 500 000 SPH particles. Comparisons were carried out for all energy and angular momentum quantities (see figures 8 and 9 in P12), separation evolution (their figure 4) and were found to be in all cases within 10 per cent of one another. Small differences can be ascribed to different resolution, differences in the codes and to a somewhat different stellar setup (more below). We carried out a convergence test using 3 resolutions. In Fig. 3.15 we show the evolution of the orbital separation for simulations using 23, 000, 230, 000, and 2.3 million particles, respectively. The factor of 10 difference between the resolutions is just larger than the minimum resolution step needed for such a test. While this test shows that we have not yet achieved formal convergence, the change in orbital evolution with resolution is much smaller between the higher two resolutions than between the lower two, indicating converging behaviour.

The final separation we obtain with PHANTOM is $21 R_{\odot}$ at ~ 180 days (the end of the dynamical in-spiral as defined above), $16 R_{\odot}$ at 1000 days and $14 R_{\odot}$ at the end of the simulation at 1500

²The softening length in PHANTOM is equivalent to the smoothing length in ENZO. PHANTOM reserves the term “smoothing length” for the size of the smoothing kernel, such that each SPH particle has a smoothing length.




FIGURE 3.15: Evolution of the separation, a , between the two particles representing the core of the primary and the companion, used to show the convergence for the PHANTOM code. The simulation reproduces the one from P12 with the same companion’s mass as this work ($M_2 = 0.6 M_\odot$). The number of SPH particles used is: 2.3×10^4 (blue), 2.3×10^5 (red), 2.3×10^6 (yellow). The inset shows a $10\times$ zoom on the end of the rapid in-spiral phase.

days. The first two values can be compared to $21 R_\odot$ at the end of the in-spiral and $18 R_\odot$ at 1000 days for simulation SPH2 of P12. The mass unbound in our PHANTOM simulation is ~ 12 per cent of the envelope mass, compared to approximately 10 per cent for SPH12 of P12. We think that these differences are mainly due to the differences in resolution and the slightly different initial separation of $100 R_\odot$ that we had to adopt because the relaxed star in PHANTOM has a larger radius ($R = 93 R_\odot$; defined using the volume-equivalent definition of Nandez et al. 2014) compared to the radius of the star stabilised in simulation SPH2 of P12 ($R = 83 R_\odot$).

We then carried out a second PHANTOM simulation with a larger initial separation, to investigate whether a larger initial separation leads to a wider final separation. Our new PHANTOM simulation has an initial separation of $218 R_\odot$, the distance at which the primary fills its Roche lobe. Ideally we would have used a larger separation of $300 R_\odot$, like for the ENZO simulation discussed here. However, the orbital evolution of a PHANTOM simulation with an initial separation of $300 R_\odot$ was too slow to reach the common envelope phase in reasonable computational times. This is due to the stability of SPH simulations to surface deformations (Springel (2010); see also our discussion in Section 3.3.3 and Fig. 3.10).

The final separation of this simulation was $29 R_\odot$, using the criterion above or $22 R_\odot$ at the end of the simulation (5050 days). In Figure 3.16 we compare the evolution of the separations




FIGURE 3.16: Evolution of the separation, a , between the two particles representing the core of the primary and the companion for the PHANTOM simulations with initial separations of $100 R_{\odot}$ (blue curve) and $218 R_{\odot}$ (red curve). For a clearer comparison of the final separation the blue line has been shifted forward in time by 3096 days, which is the time when the orbital separation of that simulation reaches $100 R_{\odot}$.

of our two PHANTOM simulations by shifting the simulation starting at $100 R_{\odot}$ by 3096 days to a time when the other simulation, starting at $218 R_{\odot}$, has a separation of $100 R_{\odot}$. We can therefore compare compare 29 and $22 R_{\odot}$ for the “wide” simulation, with $21 R_{\odot}$ and $16 R_{\odot}$, for the PHANTOM simulation starting with a smaller separation.

The two PHANTOM simulations show an increase in the final orbital separation by 38 per cent, corroborating the conclusion drawn from comparing the two ENZO simulations that the final orbital separation increases by including phases before the rapid in-spiral. The mass unbound in the PHANTOM simulation starting at $218 R_{\odot}$ is 15 per cent, marginally larger than the 12 per cent for the same simulation starting with a lower initial separation.

3.5 Comparison with published simulations

Here we carry out a comparison of CE simulations containing at least one giant (Rasio & Livio 1996, Sandquist et al. 1998, P12: Passy et al. 2012, Ricker & Taam 2012, Nandez et al. 2015, Ohlmann et al. 2016), highlighting possible trends or aspects that need further clarification. We do not include those simulations carried out by Staff et al. (2016a) that started with highly eccentric orbits. All the final results of these simulations are summarised in Table 3.1. All the simulations,

except that of Nandez et al. (2015) are carried out with codes that include similar physics and can be more directly compared.

3.5.1 Side-by-side code comparison

The only side-by-side code comparison that can be carried out is between ENZO, SNSPH and PHANTOM for which almost identical simulations were carried out. The comparison between the first two was already carried out by P12. Here we only add that SNSPH final separations are approximately 10 per cent larger than for ENZO. The relative difference does however increase for simulations with very low mass companions ($0.1 M_{\odot}$).

The comparison between SNSPH (simulation SPH2 in P12) and our own PHANTOM simulation shows that, at the “criterion point”, the two values are the same within one solar radius, while at 1000 days the PHANTOM separation is 10 per cent smaller, but has the same value as the ENZO simulation. We conclude that code-to-code differences for these three codes and for this parameter space are within 10 per cent for simulations with companions more massive than $\sim 0.3 M_{\odot}$.

3.5.2 The final orbital separation as a function of M_2/M_1

Comparing the 5 ENZO simulations of P12 with each other, or their 5 SNSPH simulations with each other, or, to an extent, comparing two of the simulations of Sandquist et al. (1998) for which only M_2 was changed, we conclude that the final separation has a clear relationship to the value of M_2 for the same value of M_1 . It is difficult to compare with the other simulations, because although two simulations may have the same value of q , the binding energies of the primaries’ envelope could be vastly different (but see Sec. 3.5.3).

Sandquist et al. (1998) also compared two simulations with different primaries and the same q . The simulation with the more extended, lower binding envelope energy primary has a much larger final separation (see Table 3.1), but we did not plot it because the final separation cannot decrease much below the resolution times the particles’ smoothing length and in that simulation the two values are almost the same.

The post-CE binary observations of Zorotovic et al. (2011) show that post-CE binaries with

post-RGB primaries (identified by a mass smaller than $0.5 M_{\odot}$) have systematically smaller separations than post-CE binaries with post-AGB primaries (which have masses larger than $0.5 M_{\odot}$). They also show a marginal correlation, though statistically “real”, between secondary mass and post-CE orbital separation. The latter conclusion is in line with the simulations, though clearly the signal in the data is diluted by the range in primary masses for each secondary mass (see below).

3.5.3 The final orbital separation as a function of primary mass or envelope binding energy

The simulations of Rasio & Livio (1996), Nandez et al. (2015) and Ohlmann et al. (2016) produce distinctly lower separations, even accounting for their different values of M_2 . We ascribe this difference to heavier and/or more compact primaries, resulting in envelopes with larger binding energies. The P12 and Sandquist et al. (1998) simulations with the most comparable values of q are extremely similar, despite the fact that the lower mass for the former should promote a wider separation than the higher mass for the latter. On the other hand, Sandquist et al. (1998) simulated a more extended AGB star, which could lead to a wider separation, countering the effect of the larger primary mass.

We do not think that the reason for the compact final configuration achieved by Nandez et al. (2015) is the extra energy source. If anything, that should have contributed to a wider separation, because of a more prompt envelope ejection. The reason is likely the more compact configuration of their RGB giant.

3.5.4 The final orbital separation as a function of giant spin at the time of Roche-lobe overflow

It could be argued that starting with a wider initial separation has, primarily, the effect of spinning up the giant, by injecting the the orbital angular momentum into the envelope. The farther the initial separation (within the limits of tidal effectiveness) the more angular momentum is available. This may in turn reduce the velocity contrast between the companion and the envelope and result in a smaller gravitational drag. However, the rotating and non-rotating simulations of Sandquist et al.

(1998) reached the same final separation, indicating that the larger amount of angular momentum of their rotating star does not influence the outcome of those CE simulations.

The simulations of Sandquist et al. (1998) did not stabilise the star in the rotating frame. Neither did those of Ohlmann et al. (2016). On the other hand those of Rasio & Livio (1996) stabilised their rotating giant in the co-moving frame while slowly decreasing the orbital separation to the moment of Roche-lobe overflow. Considering the lack of a difference between the rotating and non-rotating results of Sandquist et al. (1998) we suggest that the lack of stabilisation does not have an effect on the results.

If not by the presence of induced rotation, how can we explain the slower plunge and the increase in the final orbital separation by between 25 (Section 3.3.1) and 38 per cent (Section 3.4) that resulted from larger initial separations? Another cause for this difference could be the envelope being more extended as an effect of simulating the entire Roche-lobe overflow phase.

3.5.5 Unbound mass

The mass unbound at the end of the simulations listed in Table 3.1 ranges between 8 and 46 per cent (not counting the result of Nandez et al. (2015)), something that cannot be accounting for the fact that not all values were obtained with the same definition of bound mass.

By looking at the outputs of the simulations of Sandquist et al. (1998) and P12, one could deduce that lower mass ratios (M_2/M_1) lead to less unbound mass. However our work, that of Rasio & Livio (1996), of Ricker & Taam (2012) and of P12 show unbound gas masses that are overall lower than for the simulations of Sandquist et al. (1998) or the simulations of Nandez et al. (2015) not including recombination energy, which unbound 50 per cent of the envelope (although this is only stated in the text of that paper and no plots, nor other data are presented for that simulation).

The simulations of Staff et al. (2016a) with a $3 M_\odot$ AGB star in a common envelope with 0.6- $3.0 M_\odot$ companions have not been included in Table 3.1 because of their high initial eccentricity, which makes them stand on their own. We note, however, that resolution tests carried out in the context of those simulations show that slightly unconverged simulations tend to unbind significantly more mass than better converged simulations. We therefore wonder whether convergence,

which is seldom formally achieved in this type of time-consuming simulations, may impact the value of the unbound mass.

The impact of the recombination energy on the unbound mass was shown to be a promising avenue for further study by Nandez et al. (2015), who derived unbound masses of almost 100 per cent, compared to 50 per cent not including recombination energy. The simulations by Ohlmann et al. (2016) repeated with recombination energy (Ohlmann, private communication) also unbind twice the mass as those carried out without recombination energy and presented in Ohlmann et al. (2016), but even so, still only unbinding a small fraction of the total envelope.

3.6 Conclusions

In this chapter we have expanded on the results of P12: Passy et al. (2012) by repeating one of their common envelope simulations, a $0.88 M_{\odot}$, RGB primary and a $0.6 M_{\odot}$ companion, but increasing the initial orbital separation from 1 to 3 times the initial stellar radius. This is the approximate value of the orbital separation for which a tidal capture can be expected and as such it is the approximate value of the maximum angular momentum that can be delivered to the primary for such a system. We have also carried out a parallel set of simulations using the SPH code PHANTOM aimed at continuing code-code comparison while checking the conclusions obtained using the grid code.

We divided the evolution into a pre-contact phase, a mass transfer phase and a rapid in-spiral phase. The pre-contact phase is driven by tides. However, this phase is unrealistically short in our simulation, due to small but tidally significant oscillations of the primary star envelope set in motion by the introduction of the companion in the computational domain. Despite this unrealistic time-scale, a reasonable conservation of angular momentum ensures that the primary absorbs the correct amount of angular momentum. The mass transfer and the rapid in-spiral phases are in approximate agreement with the theoretical expectations. Starting with a larger initial separation results in a larger final separation by between 25 and 38 per cent for the set of parameters tested in this work. Based on a comparison with simulations in the literature, we conclude preliminarily that this has to be due primarily to the stellar expansion during the Roche-lobe overflow phase, rather than the extra angular momentum injected into the primary.

We observed that the unbinding of the mass happens in a short, bursting event which begins

shortly before the rapid in-spiral phase and peaks early during it, as expected from previous work. All the unbound mass is then rapidly pushed out of the simulation box. The total amount of mass unbound during the interaction is 16 per cent of the total envelope mass, while in the equivalent simulation of P12: Passy et al. (2012), 11 per cent of the envelope mass is unbound. The companion could thus eject 60 per cent more mass than for a simulation starting with a smaller orbit, probably because by tapping the reservoir of orbital angular momentum in the wider orbit the envelope has a lower binding energy. This increase is echoed by comparing our two SPH simulations that started at different initial separations.

A study on the gravitational drag has been carried out, estimating all the relevant quantities in the surroundings of the companion during the simulation. We do not find any change of the Mach regime, differently to what was observed by Staff et al. (2016b). Instead the end of the in-spiral seems to be caused mainly by a decrease in density during the later stages of the interaction. Finally, since the final separation obtained is approximately three times larger than the smoothing length we do not think that the result is affected by the approximated gravity.

By setting our results in the context of previous work, a new picture seems to be emerging, indicating that the discrepancy between observed post-CE separations and simulation is not as definitive as when Passy et al. (2012) carried out their comparison, with several simulations reproducing very small final separations, even for relatively large values of the M_2/M_1 ratio. The strong dependence of final separation on secondary mass can only be assessed by the P12 simulations, which carried out the necessary comparison. The amount of unbound mass seems to cluster in two groups, with low ($\lesssim 15$ per cent) and high ($\gtrsim 40$ per cent) values, but the reason for this difference is not clear.

Nandez et al. (2015) report to have resolved the problem of unbinding the CE by including recombination energy in their simulations. However, there is an indication that in the different regime of giant CE simulations, the inclusion of recombination energy is not sufficient to unbind the envelope (Ohlmann, private communication). It is also unclear whether the entire recombination energy budget can also be used to unbind the envelope (as would be the case in all adiabatic simulations) or whether it would partly radiate away, decreasing its efficiency. This said, the pioneering work of Nandez et al. (2015) (see also Ivanova et al. (2015)) constitutes a step that will have to be considered and tested further in future numerical simulations.

It is hoped that future simulations by different groups will attempt to clarify some of the questions above by carrying out similar simulations with a range of parameters. In this paper we have also compared the simulations with the observations previously used by P12. However additional observations, such as those by Zorotovic et al. (2011) show new trends, which can guide parameter choices of future simulations.

The results just discussed will have, however, to be partially reconsidered in future work. In fact, days before the submission of this thesis, we obtained results from an additional simulation carried out as in P12, but with the new AMR setting used here. The final separation of this (P12-AMR) simulation is $10 R_{\odot}$ instead of the $16 R_{\odot}$ obtained by P12 with ENZO and the $18 R_{\odot}$ with SNSPH, while the mass unbound in the P12-AMR simulation is 14% of the total envelope mass, instead of the 10% obtained by P12 with SNSPH. Hence our original conclusion, that including the phases prior to the rapid in-spiral results in a larger final orbital separation, is still valid. The mass of the envelope unbound becomes instead more similar to the results of P12, making debatable the conclusion that we are able to eject more mass by starting the simulation at a larger initial orbital separation. Moreover, we can consider how this result changes the ENZO vs. PHANTOM comparison. The results obtained by the comparison between the two PHANTOM simulations, starting respectively at $218 R_{\odot}$ and $93 R_{\odot}$, are still valid. In fact the P12-AMR simulation still shows a smaller separation with respect to the the ENZO AMR simulation started at $300 R_{\odot}$, trend corroborated by PHANTOM. The final values for the envelope mass unbound are much closer to what was obtained comparing the two PHANTOM simulations: just a few percent more mass is unbound when starting the simulation at a larger separation.

We think, however, that the resolution used may have a non-negligible effect on the outcomes of the simulations and therefore even run-to-run differences have to be carefully considered when evaluating the outputs of two simulations. We leave it to future work to compare resolutions, especially in the inner regions of the primary star, and explain this effect.

4

The effect of the primary rotation on the common envelope simulations

In this chapter we investigate the rotation profiles of RGB stars. The ultimate aim is to carry out the same simulation as P12, hereafter P12, including a primary star with a realistic rotation profile. The motivation for this analysis is that in previous numerical work rotation of the primary was seldom included. When this was done, it was either by spinning the star without stabilising it (grid simulations) or by setting the system in co-rotation (SPH simulations). The results of the only existing side-by-side comparison indicate that there is no difference (Sandquist et al. 1998), but few details were given and doubts remain. Even if rotation was presumed not to affect the CE from an energetic point of view (P12), a rotating star will be more extended and less dense, therefore a companion plunging into its envelope could show a different behaviour compared with a similar non-rotating model. We here aim to understand in detail the impact of rotation on CE interaction.

A second motivation is understanding the results from Chapter 3. There, the wider orbital separation efficiently spun-up the primary while also allowing it to relax and naturally overflow its Roche lobe before the rapid in-spiral. It was therefore impossible, in that chapter, to distinguish between the effects of a larger angular momentum budget and those of a different stellar geometry at the time of the in-spiral.

Since a lot of simulations with and without rotation and starting at various orbital separations are present in the literature, we decided to separate, at least for one simulation, the effect of rotation from those of a larger orbital separation. Finally, we also wanted to investigate the possibility of pre-calculating a giant star with the correct rotation profile. As it happens, we were not successful in this last attempt. To perform the latter investigation we carried out a study of giant stars rotation profiles, with which we start the chapter.

4.1 The rotation of red giant branch stars

The rotation of giant stars has been widely studied. However, in the last decade, with the *CoRoT* (Baglin et al. 2006) and *Kepler* (Borucki et al. 2010, Koch et al. 2010) satellites, more accurate observations have become available, opening up the possibility of new studies. It appears that there are two groups of RGB stars: those with slow rotating cores, making up the majority, and those with fast rotating ones (Ceillier et al. 2012 and Goupil et al. 2012). Both Ceillier et al. (2012) and Goupil et al. (2012) performed 1D stellar evolution simulations with the intent of reproducing the observed rotation profiles. As shown in Figure 4.1, taken from Ceillier et al. (2012), they were able to reproduce most stars with high core speeds, but could not account for the slow core rotators. This limitation in the ability to reproduce this type of object possibly resides in uncertainties in the description of angular momentum transport (see also the recent work of Cantiello et al. 2014). Similar results are obtained by other studies such as those by Palacios & Brun (2007), Eggenberger et al. (2010), Beck et al. (2012), Deheuvels et al. (2012), Eggenberger et al. (2012) and Marques et al. (2012).

Independently of the core, the envelopes of RGB stars rotate approximately as solid-bodies, due to the effectiveness of the transport of angular momentum. Surface rotations of giants are observed to be $1 - 2 \text{ km s}^{-1}$ (Massarotti et al. 2008), with a few cases rotating slightly faster.




FIGURE 4.1: From Ceillier et al. (2012). Rotation profile for the numerical model (solid line) failing to reproduce the observed rotation of the giant star KIC 7341231 with a slow-rotating core (top dashed line) and solid-body surface rotation (bottom dashed line).

Occasionally surface rotation velocities between 10 and 20 km s^{-1} are measured (Massarotti et al. 2008). These cases are thought to be due to spinning-up of the envelope at the hand of a binary companion that exerts a tidal force on the giant, because attempts to model fast rotation rates in single giants failed (Nordhaus & Blackman 2006, García-Segura et al. 2014). Most binaries that enter a CE phase will have been interacting tidally and it is therefore likely that most of the giant stars that are about to interact with a companion are spinning relatively rapidly. Since the giant core is not resolved in ENZO, it is only the solid-body rotating envelope that needs to be reproduced. We would therefore like a MESA model with a realistic, relatively high ($10 - 20 \text{ km s}^{-1}$) rotation profile.

4.2 One dimensional simulations with rotation

In this section we will discuss the procedure we followed in attempting to create a model of a giant with rotation, in line with the observations, as discussed in Section 4.1. Ideally, the model would

have a fast surface velocity, between 10 km s^{-1} and 20 km s^{-1} , as expected of a giant that has interacted tidally with its companion.

Pre-emptying the results of the following sections, we try various attempts to obtain a rotating star such as the one introduced above, however we do not achieve the required surface velocities. Our 1D simulations are all carried out with different parameters, but always considering a single star, without the influence of a binary companion. We took this approach because the option to perform a stellar evolution simulation with the influence of a binary companion was not available in MESA at the time we performed the computations described in this chapter. Ultimately the way forward is therefore to calculate MESA models in “binary” mode (Paxton et al. 2015), using the latest MESA binary module. However, even succeeding in calculating a MESA model with the desired rotation rates, the challenge will be in mapping such a 1D model into the 3D domain, particularly in the knowledge that stellar structure is likely non spherical.

4.2.1 Generalities of the simulations

We aim to reproduce a stellar model similar to that used by P12, but including rotation. The star used in that paper was a $1 M_{\odot}$ star, evolved to the RGB until the core reached a mass of $\simeq 0.392 M_{\odot}$, at which point the radius of the star was $\simeq 83 R_{\odot}$ and its total mass was $\simeq 0.88 M_{\odot}$ (the main parameters of the 1D stellar structure simulation are shown in Table 4.1). The model

PARAMETER	VALUE
Initial star mass	$M_{\star} = 1.0 M_{\odot}$
Initial H mass fraction	$X_{\star} = 0.7380$ (solar, Grevesse et al. 2010)
Initial He mass fraction	$Y_{\star} = 0.2485$ (solar, Grevesse et al. 2010)
Initial heavy elements mass fraction	$Z_{\star} = 0.0134$ (solar, Grevesse et al. 2010)
Reimers RGB mass-loss scheme	$\eta_{\star} = 0.5$

TABLE 4.1: P12 star parameters.

has a solid-body rotation on the zero age main sequence (ZAMS). The treatment of magnetic braking in MESA (Paxton et al. 2010, Paxton et al. 2011, Paxton et al. 2013, Paxton et al. 2015) will successfully distribute the angular momentum so as to create realistic rotation profiles.

4.2.2 Varying the ZAMS angular frequency

To start, we have varied the initial ZAMS solid-body velocity parametrized by $\omega/\omega_{\text{crit}}$, where $\omega_{\text{crit}} = \sqrt{GM/R^3}$ is the angular frequency for which the modulus of the centrifugal force becomes equal to the modulus of the gravitational attraction and the star breaks up. Here G is the gravitational constant, M is the mass of the star and R is its radius. By increasing the initial solid-body rotation velocity we expect that despite the action of magnetic braking the velocity during the giant phase will also be larger. The core preserves a significant part of the original angular momentum, which is later given to the envelope of the giant through the interaction between the two regions.

Ten simulations were carried out and stopped at the tip of RGB. All the simulated stars are in the mass range $0.5 M_{\odot} < M_{\star} < 2.25 M_{\odot}$, for which we expect the helium-flash to happen. Results are shown in Table 4.2. The simulations with $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}} \geq 0.6$ show a code instability, causing

$[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$	Age (yr)	$[m/M_{\star}]_{\text{SB}}$	$R (R_{\odot})$	$M (M_{\odot})$	$v_{\text{rot}} (\text{km s}^{-1})$	$\omega (\text{Hz})$
0.1	1.21×10^{10}	0.6	171	0.773	3.15×10^{-2}	4.22×10^{-11}
0.2	1.24×10^{10}	0.6	170	0.775	6.52×10^{-2}	8.75×10^{-11}
0.3	1.30×10^{10}	0.6	167	0.778	1.04×10^{-1}	1.41×10^{-10}
0.4	1.37×10^{10}	0.6	168	0.779	1.42×10^{-1}	1.92×10^{-10}
0.5	1.45×10^{10}	0.6	168	0.777	1.70×10^{-1}	2.32×10^{-10}

TABLE 4.2: Results for the simulations for a range of ZAMS initial frequencies, for a $1 M_{\odot}$ solar mass star. “Age” is the age at the tip of RGB, $[m/M_{\star}]_{\text{SB}}$ is the fraction of mass in the core, R is the radius of the star at the end of the simulation, M is the mass of the star at the end of the simulation, v_{rot} is the equatorial tangential velocity on the surface of the star and ω is the angular velocity.

the time-step to become extremely small. This is probably due to the extreme physical regime induced by the very high rotation velocity. A solar model with $\omega = 0.6\omega_{\text{crit}} = 3.78 \times 10^{-4} \text{ rad s}^{-1}$ has a surface velocity of $v_{\text{rot}} = \omega R \simeq 265 \text{ km s}^{-1}$. Such large velocities do not seem to be handled very well by the code. Since increasing the resolution does not help, we limited ourselves to rotation velocities below this value.

The results obtained for the angular frequency profiles are in approximate agreement with observations and with the models described in Section 4.1. In Figure 4.2 we show the angular profile velocity just before the helium core flash. Table 4.2 demonstrates that the angular frequency




FIGURE 4.2: Three representative plots for the $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}} = 0.5$ stellar model of Table 4.2 at the end of the simulation. The evolutionary track of the star in the $L - T_{\text{eff}}$ diagram is shown on the top-left plot and the core temperature-density diagram is shown in the bottom-left plot. The angular frequency profile (labelled as “omega”, thick yellow line) along with the tangential rotation velocity (v_{rot} , thin purple line) are plotted as a function of m/M_{\star} (labelled as “q”) in the plot on the right.

of the solid-body rotating envelope increases with increasing $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$, as expected, but also shows that, even with our artificially high ZAMS rotation rates, we are not able to reach the rotation speeds characterising giants in binary systems (but see Section 4.2.4).

The results shown above help to get an idea of how the uniform angular frequency at the ZAMS could act as a usable parameter to manage our models, but they are far away from the large surface rotation rates we need. In fact if we look at Figure 4.2 the first thing we notice is the location of the flattening of the angular frequency, at approximately $m/M_{\star} = 0.6$ in every one of the final models. We call this parameter $[m/M_{\star}]_{\text{SB}}$, for solid body in all tables. Let us consider now the ratio $m_{\text{core}}/M_{\star}$ (where m_{core} is the mass of the core of the star), which represents the percentage of the total mass residing in the core; in P12 their star has been evolved until $m_{\text{core}}/M_{\star} \simeq 0.445$, thus a star taken at the moment prior to the helium flash is too evolved for our purpose.

There are three paths that we can follow to compare one of our rotating stars with the one of P12. We can consider the rotating MESA simulation when the core of the RGB star is $m_{\text{core}} \simeq 0.392 M_{\odot}$, when $M_{\star} \simeq 0.88 M_{\odot}$ or when $R_{\star} \simeq 83 R_{\odot}$. The data obtained from the same set

of simulations shown in Table 4.2, but stopped at the evolutionary stage where the core mass is $\simeq 0.392 M_{\odot}$ are shown in Table 4.3. What appears clear is that the presence of rotation expands the

$[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$	Age (yr)	$[m/M_{\star}]_{\text{SB}}$	$R (R_{\odot})$	$M (M_{\odot})$	$v_{\text{rot}} (\text{km s}^{-1})$	$\omega (\text{Hz})$
0.1	1.21×10^{10}	0.436	92	0.9	2.96×10^{-1}	7.38×10^{-10}
0.2	1.24×10^{10}	0.437	92	0.9	5.90×10^{-1}	1.47×10^{-9}
0.3	1.30×10^{10}	0.437	92	0.9	8.78×10^{-1}	2.16×10^{-9}
0.4	1.37×10^{10}	0.438	92	0.9	1.15	2.86×10^{-9}
0.5	1.48×10^{10}	0.438	92	0.9	1.41	3.47×10^{-9}

TABLE 4.3: Results for the same simulations of Table 4.2, stopped at the evolutionary stage when $m_{\text{core}} \simeq 0.392 M_{\odot}$.

star marginally with respect to the model used in P12 (as one would expect), but also accelerates its evolution. In fact one can see that the total mass for every model is slightly greater than the one of the reference paper at the same core mass, where less mass loss is expected due to the absence of rotation. Said in another way: when the core has burned the same amount of material the rotating stars have lost less mass than the reference, although one could expect them to lose their mass more rapidly because of rotation.

Overall, the models obtained in this way are very similar to the reference one, but the surface rotation velocities achievable in this way cannot exceed $\simeq 2 \text{ km s}^{-1}$. If an increase of the range of $[v_{\text{rot}}]_{\text{SURF}}$ is needed, one must modify other parameters of the initial model.

The data obtained from a set of simulations like those presented in Table 4.2, but stopped at the evolutionary stage where the total mass is $\simeq 0.88 M_{\odot}$ are shown in Table 4.4. This criterion

$[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$	Age (yr)	$[m/M_{\star}]_{\text{SB}}$	$R (R_{\odot})$	$M (M_{\odot})$	$v_{\text{rot}} (\text{km s}^{-1})$	$\omega (\text{Hz})$
0.1	1.21×10^{10}	0.470	110	0.88	1.92×10^{-1}	3.96×10^{-10}
0.2	1.25×10^{10}	0.470	110	0.88	5.00×10^{-1}	7.96×10^{-10}
0.3	1.30×10^{10}	0.469	110	0.88	7.50×10^{-1}	1.19×10^{-9}
0.4	1.37×10^{10}	0.469	110	0.88	7.68×10^{-1}	1.59×10^{-9}
0.5	1.48×10^{10}	0.468	110	0.88	9.61×10^{-1}	2×10^{-9}

TABLE 4.4: Results for the same simulations of Table 4.2, stopped at the evolutionary stage when $M_{\star} \simeq 0.88 M_{\odot}$.

aims to create a model with the same total mass of the star in P12. As expected, in all the models for this selection criterion the total radius of the star is greater than the radius of the star in the reference paper by about $27 R_{\odot}$. For this set of models the surface rotation velocities cannot exceed $\simeq 1 \text{ km s}^{-1}$.

The data obtained from a set of simulations like those presented in Table 4.2, but stopped at the evolutionary stage where the total radius is $\simeq 83 R_{\odot}$ are shown in Table 4.5. For all models

$[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$	Age (yr)	$[m/M_{\star}]_{\text{SB}}$	$R (R_{\odot})$	$M (M_{\odot})$	$v_{\text{rot}} (\text{km s}^{-1})$	$\omega (\text{Hz})$
0.1	1.22×10^{10}	0.421	83	0.92	3.60×10^{-1}	9.93×10^{-10}
0.2	1.24×10^{10}	0.421	83	0.92	7.15×10^{-1}	1.97×10^{-9}
0.3	1.30×10^{10}	0.421	83	0.92	1.07	2.94×10^{-9}
0.4	1.36×10^{10}	0.421	83	0.92	1.42	3.9×10^{-9}
0.5	1.44×10^{10}	0.422	83	0.92	1.74	4.79×10^{-9}

TABLE 4.5: Results for the same simulations of Table 4.2 at the evolutionary stage when $R_{\star} \simeq 83 R_{\odot}$.

obtained with this criterion the total mass when $R_{\star} \simeq 83 R_{\odot}$ is greater than the one of P12 ($M_{\star} = 0.88 M_{\odot}$). This is probably due to the effect of the centrifugal force. In fact this force gives rise to two effects simultaneously: the mass loss and the expansion of the star. Now, the balance of these two phenomena is what determined the increased mass at similar radius for these models. Although there is mass loss, the effect of the centrifugal force on the expansion of the star is stronger, hence the star reaches the desired radius before the total mass decreases to $0.88 M_{\odot}$. For this set of models the surface rotation velocities cannot exceed $\simeq 2 \text{ km s}^{-1}$. None of our models achieve sufficient surface rotation.

4.2.3 Varying η_{\star} and M_{\star}

Here we try and increase the angular frequencies by varying the Reimers mass-loss parameter, η_{\star} , and the main sequence mass, M_{\star} . Three simulations have been carried out to check the effects of varying the Reimers RGB wind parameter on the final angular frequency of the star, one for each of the three criteria, all with the same $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$ of 0.5, the one that yielded the highest angular frequencies in every case. We chose a very low value for η_{\star} of 0.0001, which means almost no

mass loss, since the higher the mass loss, the more the angular momentum loss. The results are shown in Table 4.6. Clearly, the low value of η_* does not significantly help us to spin-up the star

$[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$	Age (yr)	$[m/M_*]_{\text{SB}}$	$R (R_\odot)$	$M (M_\odot)$	$v_{\text{rot}} (\text{km s}^{-1})$	$\omega (\text{Hz})$
0.5	1.53×10^{10}	0.396	92	1	2.91	7.24×10^{-9}
0.5	1.55×10^{10}	0.625	206	<u>0.88</u>	3.73×10^{-1}	4.14×10^{-10}
0.5	1.47×10^{10}	0.421	<u>83</u>	0.92	1.75	4.82×10^{-9}

TABLE 4.6: Results for three simulations with $\eta_* = 10^{-4}$ and $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}} = 0.5$. The parameter that matches the equivalent P12 is underlined. For the first model in this table the parameter is $m_{\text{core}} \simeq 0.392$.

to the rate expected for a giant that is tidally interacting with a companion. At most we reach $\simeq 3 \text{ km s}^{-1}$ for a model with the same core mass as P12.

Three additional simulations were carried out to check the effect of varying the total mass of the star on its final angular frequency, one for each criterion, but with the same $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$ of 0.5, the one that yielded the highest angular frequencies in every case. We chose a value for M_* of $1.3 M_\odot$. Since varying the mass can significantly change the evolution of the star we did not want to deviate too much from the model of P12. We expect that increasing the total initial mass, at fixed $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$, more angular momentum is imparted to the model. The results are shown in Table 4.7. Overall the results show that increasing the mass of the model does not achieve

$[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$	Age (yr)	$[m/M_*]_{\text{SB}}$	$R (R_\odot)$	$M (M_\odot)$	$v_{\text{rot}} (\text{km s}^{-1})$	$\omega (\text{Hz})$
0.5	5.08×10^9	0.322	86	1.23	1.43	3.77×10^{-9}
0.5	5.22×10^9	0.999	287×10^2	<u>0.88</u>	2.47×10^{-2}	1.97×10^{-11}
0.5	5.08×10^9	0.318	<u>83</u>	1.23	1.50	4.14×10^{-9}

TABLE 4.7: Results for the three simulations with $M_* = 1.3 M_\odot$. The parameter that matches the equivalent P12 is underlined. For the first model in this table the parameter is $m_{\text{core}} \simeq 0.392$.

sufficiently large rotation velocities.

4.2.4 A work-around of code limitations for fast rotating stars

With an overall maximum surface rotation velocity of 2 km s^{-1} for the standard models of Table 4.5 and of 3 km s^{-1} for the modified models of Table 4.6, the star has clearly not been spun up

sufficiently so as to reach the range of velocities expected from observations of giants in binaries (Massarotti et al. 2008, $10 - 20 \text{ km s}^{-1}$). An option to simulate a CE with a spinning star could be to take the model in a lower part of the RGB branch, which will have higher rotation velocities.

A way to perform simulations with $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}} > 0.5$ was found by considering further the approximations MESA has to make to model rotating stars¹. Although this methodology did not result in a viable rotating model, we outline it here as part of a corpus of knowledge on coupling 1D and 3D models to understand binary interactions.

Stellar rotation in MESA is implemented following the prescription of the classic paper of Endal & Sofia (1976), which applies corrections to the equations of the stellar structure due to the rotation (f_T and f_P). This model considers equipotential surfaces in the interior of a star, denoted by ψ , used as reference point to model the stellar structure. General quantities related to the equipotential surfaces of the rotating star that do not depend on the angle are:

- r_ψ , the radius of a sphere enclosing the volume $V_\psi = \frac{4}{3}\pi r_\psi^3$, that is the volume of a given equipotential surface ψ .
- M_ψ , the mass interior to the equipotential surface ψ .
- S_ψ , the surface area of an equipotential surface ψ .
- $\langle g \rangle$, the local effective gravity averaged over the polar angle θ ($0 < \theta < \pi$). Note that the star is considered cylindrically symmetric.
- $\langle g^{-1} \rangle$, the reciprocal of the local effective gravity averaged over the angle θ .

The modified equations of the stellar structure are:

$$\frac{\partial P}{\partial M_\psi} = \frac{-GM_\psi}{4\pi r_\psi^4} f_P, \quad (4.1)$$

$$\frac{\partial \ln T}{\partial \ln P} = \min \left[\nabla_{\text{ad}}, \nabla_{\text{rad}} \frac{f_T}{f_P} \right], \quad (4.2)$$

with

$$f_P = \frac{4\pi r_\psi^4}{GM_\psi S_\psi \langle g^{-1} \rangle}, \quad (4.3)$$

¹The suggestion was obtained by consulting the MESA user community.

$$f_T = \left(\frac{4\pi r_\psi^2}{S_\psi} \right)^2 \frac{1}{\langle g \rangle \langle g^{-1} \rangle} . \quad (4.4)$$

A spherical stellar structure is computed as a function of r_ψ , which approximates the structure of the rotating, non-spherical star by using the average quantities and the global properties of the equipotential surfaces of the rotating star. Once the model is computed for the average quantities, it assumes that they are valid also for the rotating non-spherical star. With this assumption the approximate shape of the equipotential surfaces can be obtained along with other angle-dependent quantities, such as the radii at the equator and pole. This approximation is valid only for small distortions and does not include the direct effect of rotation on convection (e.g., rotation could inhibit convection in some zones of the star).

When rotation is turned on, MESA starts using the corrected stellar structure equations (Equations 4.1 and 4.2). The ratio $\omega/\omega_{\text{crit}}$ at the equator is constant. This means that MESA's outputs are inherently along the equator. Since the model is valid just for small distortions, if the star is spinning too fast the corrections f_P and f_T become too large for the model to converge.

To let MESA evolve correctly an initially fast-rotating star, one must consider the limits of the Endal & Sofia (1976) model. The optimal lower values for the corrections f_P and f_T are $f_{P,\text{min}} = 0.75$ and $f_{T,\text{min}} = 0.95$, with $0 < f_P, f_T < 1$ (where the value of unity means no rotation). If the computed f_P and f_T are lower due to large imposed rotations they are reset respectively to 0.75 and 0.95. In such a way MESA can be forced to evolve stars with initial angular frequencies $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}} > 0.5$. The price of this modification is an inconsistency in the model on the ZAMS between Equations 4.1-4.2 and 4.3-4.4. During the RGB, where the star slows down significantly because of expansion, magnetic braking and angular frequencies are within the limits of the approximation and no further resetting of the values of f_P and f_T takes place.

Here we summarize the results of the simulations carried out with values of $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}} > 0.5$, using the usual three stopping criteria shown in the previous sections. For models stopped when $m_{\text{core}} \simeq 0.392 M_\odot$ results are shown in Table 4.8. These results show an increasing of the surface rotation velocity up to 1.82 km s^{-1} , decreasing for increasing $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$, probably due to the high mass-loss rates because of the high initial angular frequency. The resulting loss of angular momentum is enough to spin down the star.

For models stopped when $M_{\text{tot}} \simeq 0.88 M_\odot$ results are shown in Table 4.9. For the last two

$[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$	Age (yr)	$[m/M_{\star}]_{\text{SB}}$	$R (R_{\odot})$	$M (M_{\odot})$	$v_{\text{rot}} (\text{km s}^{-1})$	$\omega (\text{Hz})$
0.6	1.64×10^{10}	0.439	94	0.9	1.65	4.01×10^{-9}
0.7	1.82×10^{10}	0.441	95	0.9	1.82	4.39×10^{-9}
0.8	2.51×10^{10}	0.570	104	0.7	5.13×10^{-1}	1.13×10^{-9}
0.9	4.52×10^{10}	0.846	101	0.47	1.48×10^{-1}	3.33×10^{-11}

TABLE 4.8: Results for the simulations with $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}} > 0.5$, stopped when $m_{\text{core}} \simeq 0.392 M_{\odot}$.

$[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$	Age (yr)	$[m/M_{\star}]_{\text{SB}}$	$R (R_{\odot})$	$M (M_{\odot})$	$v_{\text{rot}} (\text{km s}^{-1})$	$\omega (\text{Hz})$
0.6	1.60×10^{10}	0.467	109	0.88	1.15	2.4×10^{-9}
0.7	1.82×10^{10}	0.465	109	0.88	1.33	2.8×10^{-9}
0.8	1.75×10^{10}	0.379	0.90	0.88	173	4.38×10^{-5}
0.9	1.44×10^{10}	1.23×10^{-7}	0.61	0.88	139	5.20×10^{-5}

TABLE 4.9: Results for the simulations with $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}} > 0.5$, stopped when $M_{\star} \simeq 0.88 M_{\odot}$.

models the mass limit is reached while the star is still on the main sequence, as one can see looking at the radius values. The last two models cannot be considered giants and can be removed from our sample.

Finally, for models stopped when $R_{\text{tot}} \simeq 83 R_{\odot}$ results are shown in Table 4.10. In this case a

$[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$	Age (yr)	$[m/M_{\star}]_{\text{SB}}$	$R (R_{\odot})$	$M (M_{\odot})$	$v_{\text{rot}} (\text{km s}^{-1})$	$\omega (\text{Hz})$
0.6	1.65×10^{10}	0.420	83	0.92	2.10	5.76×10^{-9}
0.7	1.82×10^{10}	0.420	83	0.91	2.36	6.51×10^{-9}
0.8	2.51×10^{10}	0.517	83	0.73	1.03	2.85×10^{-9}
0.9	4.52×10^{10}	0.730	83	0.51	0.2	5.41×10^{-10}

TABLE 4.10: Results for the simulations with $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}} > 0.5$, stopped when $R_{\star} \simeq 83 R_{\odot}$.

similar behaviour to the first criterion is observed, with an increase of the surface rotation velocity followed by its decrease due to the loss of angular momentum.

Overall we witness an increase of the maximum value of the surface rotation velocity for every criterion considered, with respect to models calculated with lower $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$. However, this increase only brings the maximum value of the rotation velocity to $[v_{\text{rot}}]_{\text{SURF}} \simeq 2.4 \text{ km s}^{-1}$,

still well below the $\simeq 10 \text{ km s}^{-1}$ target we wanted to achieve to reproduce tidally accelerated stars. Clearly there is a limit $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}}$ for which MESA has additional problems in correctly computing the models, due to the extreme regimes. The last two stars have in fact characteristics that largely diverge from the stellar parameters obtained at $[\omega/\omega_{\text{crit}}]_{\text{ZAMS}} = 0.7$. Given that we were anyway unable to reproduce the desired surface velocities, we did not investigate the causes of this divergence in the last two models.

Below we carry out a simulation with a star that is spinning and not stabilised, as is done in most grid simulations (e.g., Sandquist et al. 1998). The aim is to obtain a side-by-side comparison with the similar, non-rotating simulation of P12.

4.3 Three dimensional simulations of the CE phase with rotation

Despite the fact that we were not able to reproduce a realistic one-dimensional profile for a single RGB star, we simulate here a binary system with a rotating giant by applying solid-body rotation to the initial stellar model of P12. Below we first describe how the rotating model was setup, then we analyse the main outputs of the simulation emphasizing the differences with the results of P12. As discussed in Section 3.6, we have carried out the same simulations as P12, but using the AMR code and setup used for the rotating star, so as to minimise differences (hereafter we will refer to the non-rotating P12 simulation, but carried out with AMR as P12-AMR).

4.3.1 Simulation setup

P12 used two different numerical codes (ENZO and SNSPH) to simulate a CE interaction between an RGB giant and a compact companion, modelled as a point mass, placed on the surface of the primary. The two stars orbit each other in a circular orbit. We use the same setup as theirs, with the following exceptions: we use an updated version of ENZO that includes AMR (Passy & Bryan 2014), as we did in Chapter 3. To carry out a better comparison between a rotating and a non-rotating star, we have repeated the P12 simulation with the AMR code and the same setup/resolution of Chapter 3.

Before comparing rotating and non-rotating simulations, we comment on our new P12-AMR simulation. The final separation achieved $10 R_{\odot}$ vs. $16 R_{\odot}$ and $18 R_{\odot}$ of the P12 grid and SPH simulations, respectively. The unbound mass is 14% of the total envelope mass, compared to 10% of the P12 SPH simulation. These differences must be down to small resolution and smoothing length differences (see Chapter 3).

In the literature simulations including rotating stars have been run following two techniques: by spinning up the primary, applying no further stabilisation (Sandquist et al. 1998, Ricker & Taam 2012, Ohlmann et al. 2016), or by setting up the binary system in a co-rotating frame and relaxing the configuration (Rasio & Livio 1996, Nandez et al. 2014, Nandez et al. 2015). The latter approach consists in applying a linear friction term to all the velocities in the system in the co-rotating frame. The binary has initially a large orbital separation and the stars are slowly brought closer artificially until the desired initial separation is reached. This procedure of reducing the orbital distance has to be done on time-scales greater than those at which the linear friction term is acting to allow the star to properly relax. We note that in the co-rotating frame both orbital and primary rotation velocities are hidden and therefore not subject to the linear friction. ENZO does not have a co-rotating frame option, hence we adopted the first approach, as it allows us a direct comparison with P12.

In a binary system the stars have evolved together since birth and we expect the two stars to have adapted to each other, reaching a stable configuration. The most stable configuration for any given binary separation is when the two components are co-rotating, meaning that the orbital frequency equals the frequency of rotation of the stars around their own axes. This situation of equilibrium can be broken due to different reasons. In our case it is likely that a Darwin instability (Darwin 1879) can change a stable, synchronised system into an unstable one, where the orbital separation is reduced until Roche-lobe overflow and ultimately a CE takes place. See Section 1.2 for a description of the Darwin instability. In this case the unstable regime is triggered and the two components tidally drag each other, reducing the orbital separation until eventually an interaction starts. This phase cannot be simulated in ENZO even if ENZO were setup to simulate a binary in the co-rotating frame, reproducing tides would require a far superior precision of the flow simulation than we currently have. Additionally, even if we could simulate tides, the simulation would have to be run for unrealistically long times.

We map a non-rotating giant evolved as in P12 and Chapter 3 into ENZO and give it solid-body rotation. Artificially forcing the model to rotate makes it expand. With an initial orbital separation such as the one P12 are using, the rapid in-spiral is triggered as soon as the simulation starts, leading to gas flows that are faster than those due to the expanding star. As a result, the lack of realism is mitigated by the presence of a factor of greater magnitude which drives the CE interaction.

We give our RGB star solid-body rotation equal to 95% of the orbital frequency. This choice is in line with previous work (Ricker & Taam 2008, Ricker & Taam 2012, Ohlmann et al. 2016). This also leaves a small difference in rotation between the gas and the companion which facilitates the onset of in-spiral. The chosen rotation corresponds to a primary's surface velocity of 33 km s^{-1} , somewhat larger, though comparable to values of giants surface velocities in binary systems (Massarotti et al. 2008). The orbital velocities with respect to the centre of mass of the system are $v_1 \simeq 23.6 \text{ km s}^{-1}$ and $v_2 \simeq 34.7 \text{ km s}^{-1}$, where v_1 is the velocity of the primary's core and v_2 is that of the companion.

4.3.2 Evolution of the binary separation

The orbital evolution of the rotating and non-rotating simulations are almost identical (Figure 4.3, upper panel). The similarity of the two simulations is somewhat unexpected. On the other hand, simply on energetic arguments the addition of rotation adds

$$E_{\text{kin,rot}} = \frac{1}{2} \left(\frac{2}{5} M_1 R_1^2 \right) \left(\frac{v_{\text{surf}}}{R_1} \right)^2 \simeq 9.53 \times 10^{44} \text{ erg}, \quad (4.5)$$

where M_1 is the mass of the primary, R_1 is the radius of the primary and v_{surf} is the surface velocity of the primary. This value is two orders of magnitude smaller than any other form of energy in the star (internal or potential). We started our simulation from a slightly smaller initial separation than P12, because by using AMR the initial RGB model has a radius of $\simeq 83 R_{\odot}$ instead of $\simeq 86 R_{\odot}$. The value of the final separation at 400 days obtained for the rotating model is plotted, together with the value of P12-AMR and previous work (Chapter 3) in Figure 4.4.

As explained in detail in Chapter 3, the main process driving the rapid in-spiral is gravitational drag, which increases with the density of the medium and the relative velocity between it and




FIGURE 4.3: Upper panel: evolution of the separation between the core of the primary and the companion for the simulation with a rotating primary (blue), for the simulation P12-AMR (cyan) and for the separation averaged over one orbital cycle (red), as defined by P12 and in Chapter 3. The black vertical line represents the end of the rapid in-spiral phase. Lower panel: evolution of the orbital decay, computed on the separation averaged over one orbital cycle.

the companion. If we follow the formulation used in Iben & Livio (1993) and P12 for subsonic regimes:

$$F_{\text{gravodrag}} \sim \pi \rho v_{\text{rel}}^2 \left(\frac{2GM_2}{v_{\text{rel}}^2 + c_s^2} \right)^2 \sin \alpha, \quad (4.6)$$

where ρ is the density of the gas around the companion, v_{rel} is the relative velocity between the companion and the envelope around it, G is the gravitational constant, M_2 is the mass of the companion, c_s is the local sound speed and α is the angle between the velocity of the companion and that of the envelope, the term inside the parentheses has the dimension of length and is known as the accretion radius (R_a). R_a is the scale length on which the gravitational drag operates, therefore all the previous quantities can be considered as their average values inside a radius R_a around the companion. In our case the initial rotation of the star reduces the relative velocity between envelope and companion. We might assume that this decreases the gravitational drag, and, at least initially, the effectiveness of the in-spiral. However, what we see is a very similar in-spiral. This is somewhat different from what was discovered by Sandquist et al. (1998), who attribute their observed larger eccentricity in the rotating simulation to the fact that when the gas has a smaller relative velocity with respect to the companion, the companion has more time to transfer its angular momentum to the gas.

The initial separation used for the calculation in Chapter 3 is $300 R_\odot$. By the time the companion reaches $83 R_\odot$, the radius of the undisturbed giant, it has deposited into the envelope a certain amount of angular momentum:

$$\Delta L_{\text{sep}} = (a_{1,i}M_1v_{1,i} + a_{2,i}M_2v_{2,i}) - (a_{1,f}M_1v_{1,f} + a_{2,f}M_2v_{2,f}) \simeq 2.15 \times 10^{52} \text{ g cm}^2 \text{ s}^{-1}. \quad (4.7)$$

On the other hand, our artificially spun up giant was imparted an angular momentum:

$$L_{\text{rot}} = \frac{2}{5} M_1 R_1^2 \left(\frac{v_{\text{surf}}}{R_1} \right) \simeq 1.33 \times 10^{52} \text{ g cm}^2 \text{ s}^{-1}, \quad (4.8)$$

where $a_{1,i}/a_{1,f}$ are the distances of the primary from the centre of mass of the system in the initial and final configurations, $a_{2,i}/a_{2,f}$ are the same for the companion, $v_{1,i}/v_{1,f}$ are the circular orbital velocities of the primary with respect to the centre of mass of the system in the initial and final configurations and $v_{2,i}/v_{2,f}$ are the same for the companion.




FIGURE 4.4: Final orbital separation vs mass ratio $q = M_2/M_1$ for observed post-CE systems (Zorotovic et al. 2010 and De Marco et al. 2011, black dots) and for simulations (Sandquist et al. 1998, green circles; Ricker & Taam 2012, cyan triangle; Rasio & Livio 1996, magenta pentagon; PHANTOM simulations, blue hexagons, see Section 3.4; Ohlmann et al. 2016, grey diamond; Nandez et al. 2015, pink cross; the results of the ENZO simulation described in Chapter 3, yellow star; simulation P12-AMR, yellow square; the results of this chapter, yellow rectangle).

These values are very similar. This tells us that the nominal rotation rates of the giants in the two numerical experiments are similar. The difference between the simulation in Chapter 3 and this one must therefore be that by the time any companion-gas interaction takes place the giant is substantially more extended and less dense for the simulation in Chapter 3, thus reducing the strength of the interaction throughout the simulation.

Finally, from Figure 4.3 (upper panel), we see a decrease in the period towards the end of both the rotating star simulation and P12-AMRs. This is not a physical effect, but rather a resonance between the orbital frequency and the code output saving frequency.

4.3.3 Envelope ejection

Only the simulation of Nandez et al. (2015) has been able to fully eject the envelope in a CE simulation by using recombination energy of the gas during the interaction. Their result, however, is discordant with the results of Ohlmann et al. (2016b) (in preparation and private communication), who, while showing a doubling of the unbound CE mass, from 8% to 16% by including recombination energy, did not succeed in ejecting the envelope.

Following the same scheme we used in Chapter 3, we define the envelope's gas in a simulation cell to be unbound when $E_{\text{tot}} = E_{\text{kin}} + E_{\text{pot}} + E_{\text{therm}} > 0$, where E_{tot} is the total energy, E_{kin} is the kinetic energy, E_{pot} is the potential energy and E_{therm} is the thermal energy. We will also discuss the contribution of the thermal energy to the total energy, as we did in Chapter 3. The envelope lifting happens mainly on the orbital plane, where energy and angular momentum are deposited, and it starts as soon as the simulation begins. The in-spiral is very quick, as shown in the previous section, and this generates a very high velocity wave of material expanding into the external medium. This wave is mainly composed of unbound gas, as shown in Figure 4.5 (first two rows), and starts leaving the simulation box at $\simeq 90$ days from the beginning of the simulation. This first unbinding event is mainly generated by gas accelerated above the escape velocity and not by heating (the difference between left and right columns in the first two rows of Figure 4.5 is small).

In the second row of Figure 4.5 it is possible to see that in the zones “behind” the expanding wave of previously unbound gas that is about to leave the domain additional unbinding, this time due to extra heating (compare left and right columns in Figure 4.5). Next, at $\simeq 263$ days (see Figure 4.5, third row), a new unbinding event is generated mainly by the heating. This happens because of successive pressure waves smashing into one another and heating the gas. Again, the unbound layer is rapidly pushed out of the simulation box. To clarify this behaviour we also plot slices of temperature and velocity in Figure 4.6 and Figure 4.7 respectively.




FIGURE 4.5: Left panel, left column: density slices perpendicular to the z axis in the orbital plane after (from top to bottom) 44, 88, 263 and 285 days from the beginning of the simulation. The point-mass particles representing the core of the primary and the companion are shown as black dots, while the white regions represent the unbound gas. The size of the black dots is not representative of any property of the point masses and is chosen only to highlight them. Left panel, right column: same as the left column, but excluding thermal energy (E_{th}) in the computation of the bound/unbound mass elements. Right panel, left and right columns: same as for the left panels, but the density slices are taken perpendicular to the orbital plane, at $x = 0$.




FIGURE 4.6: Same as the left and right panels of Figure 4.5, but the showing the temperature distribution instead.




FIGURE 4.7: Same as the left and right panels of Figure 4.5, but the showing the velocity distribution instead.

Observing the density slices in Figure 3.5, we see that the main mechanism of unbinding is via heating at the beginning of the simulation and by acceleration above escape velocity later. This is exactly the opposite of what we observe in the case of the rotating giant (Figure 4.5). Rotation favours unbinding on the outermost layers via acceleration at the beginning of the simulation. The difference can be ascribed again to the different modality of the interaction: the outer layers do not have time to adapt to the rotation when they interact with the companion, while for the work in Chapter 3 the entire star receives energy and angular momentum from the orbit while the companion is outside the envelope. Nevertheless, the amount of mass unbound here and in Chapter 3 is similar.

The evolution of the mass inside the simulation domain (Figure 4.8, upper panel) shows that no significant mass is lost from the domain until $\simeq 70$ days (green line). Mass starts to be unbound at the beginning of the simulation (cyan line), peaking at $\simeq 100$ days, but then both bound and unbound gas are pushed out of the domain as the simulation evolves. The cumulative mass leaving the domain (Figure 4.8, lower panel) better shows that gas starts to leave at $\simeq 70$ days. Initially all the mass leaving the domain is unbound, joined at $\simeq 300$ days by bound gas, which keeps increasing. Following the cumulative unbound mass line, it is possible to see the double unbinding event described earlier, with a steep initial accumulation of unbound mass, which then flattens out until $\simeq 300$ days, where it rises again because of mass leaving the box in the second unbound wave. The total unbound mass is $\simeq 15\%$ of the total envelope mass, almost identical to the P12-AMR simulation, which unbound $\simeq 14\%$ of the total envelope mass. Sandquist et al. (1998) obtain a difference in the amount of bound and unbound mass generated with a rotating and non-rotating model. Their rotating simulation unbinds $\simeq 26\%$ of the envelope mass, while their non-rotating one unbinds $\simeq 33\%$. They do not, however, discuss this difference.




FIGURE 4.8: Upper panel: evolution of the gas mass inside the simulation box over time. The blue line represents the value of the initial gas mass contained in the box and is plotted for comparison, while the green line shows the evolution of the total mass contained inside the box. The red and cyan lines show, respectively, the bound and unbound components of the mass. Lower panel: cumulative mass of the gas flowing out of the simulation box over time. Line colours have the same meaning as for the upper panel.

5

The effect of massive primaries on the common envelope interaction

Our work philosophy is to isolate specific initial parameters or conditions and determine the impact on the output of the CE simulations. This allows us, among other things, to gauge the importance of aspects of the simulations that are not included in numerical work at the moment and which might or not affect the results. Following this line of research, in this chapter we analyse the effect of a more massive RGB primary while keeping the companion the same as those used by P12, hereafter P12, to make a direct comparison. *The set of five simulations analysed in this chapter were performed by our collaborator J.-C. Passy using Canadian supercomputing resources. The outputs were then fully analysed by R. Iaconi.*

5.1 Simulations setup

The binary systems simulated in this chapter have been created in the same way as those used by P12: ENZO is used in uniform static grid mode and the companions are placed on the surface of the primary at the beginning of the simulation. The primary has been mapped and stabilised in three dimensions starting with a one-dimensional model, obtained from a stellar evolution code (in this case MESA, Chapter 2). However, here, the RGB star has an initial mass of $2 M_{\odot}$ instead of $1 M_{\odot}$. The star has been evolved until its core mass is the same as the one of P12, $M_c = 0.392 M_{\odot}$. At this point of the evolution, the star has a total mass of $M_1 = 1.97 M_{\odot}$ (compared with $0.88 M_{\odot}$ in P12) and a radius of $R_1 = 66 R_{\odot}$ (compared with $83 R_{\odot}$ in P12). This is therefore a more compact, bound star. The reason for choosing to compare two stars with equal core mass was that once the CE takes place, and assuming a binary emerges, the two binaries would be indistinguishable. The companion masses are $0.1 M_{\odot}$, $0.15 M_{\odot}$, $0.3 M_{\odot}$, $0.6 M_{\odot}$ and $0.9 M_{\odot}$, as was the case in P12. They are all meant to be compact objects, main sequence stars or white dwarfs, not resolved in the ENZO grid and therefore they are modelled as point-masses interacting only gravitationally with both gas and other particles. The primary (both the core particle and the gas) and the companion are then given Keplerian circular orbital velocities.

The simulations have been run with three different resolutions: 128^3 , 256^3 and 512^3 cells, with a static uniform grid. The size of one side of the computational domain cube is $2.4 \times 10^{13} \text{ cm} \simeq 1.60 \text{ AU} \simeq 345 R_{\odot}$ for all the three resolutions, as was the case for P12. The initial models for the primary star, after mapping them from MESA to ENZO, for the three resolutions we used are shown in Figure 5.1. With increasing ENZO resolution the MESA model is reproduced more accurately and with more sample points, especially in the core, where more MESA cells are lost due to the interpolation (see Chapter 2). After the relaxation process (again, see Chapter 2) the stellar profile slightly expands; this effect is smaller at higher resolution. This is clear from Figure 5.2, where the initial configuration of the binary system used in the simulations is shown. Note the position of the companion at the surface of the primary envelope, as was the case in P12.

P12 run simulations with 128^3 and 256^3 cells, together with SPH simulations with 500,000 particles, and concluded that the 128^3 simulations were not converged. For example, the final separation obtained with the 128^3 simulations was different from the values obtained with 256^3




FIGURE 5.1: Density profiles of the $2 M_{\odot}$ RGB star at the three ENZO resolutions used and for the original 1D MESA model.

cells and SPH, especially for low mass companions. The 256^3 cells and SPH simulations resulted instead in similar values for the final separations. This does not mean that simulations are fully converged. Here we run the 512^3 set to improve the resolution test, indicating converging behaviour. The results show that the difference between the 256^3 and the 512^3 are small, while the 128^3 show larger variations. Considering again the values for the final separations, for 0.1, 0.15, 0.3, 0.6, 0.9 M_{\odot} companions we obtain respectively 4.5, 4, 3.8, 5.9, 5.1 R_{\odot} (128^3 cells) and 2.3, 2.6, 2.2, 3.1, 4.4 R_{\odot} (256^3 cells). Also, similarly to P12, less convergence is seen at lower companion masses. The results and plots below are for the 512^3 simulations. Their initial conditions and final outputs are shown in Table 5.1, together with the results for the 256^3 and SPH simulations of P12. Note that the initial separation does not exactly correspond to the radius of 66 R_{\odot} mentioned above because of the slight increase in radius due to the relaxation process discussed above.

M_1 (M_\odot)	$M_{1,c}$ (M_\odot)	R_1 (R_\odot)	Giant	M_2 (M_\odot)	q	a_i/R_1	Ω/ω^1	Code ²	Resolution (Part./ R_\odot)	τ_{run} (day)	a_f^3 (R_\odot)	M_{Unb}^4 (%)	Ref. ⁵
1.97	0.39	69	RGB	0.1	0.05	1.0	0	u-grid(E)	0.7	250	1.3	2	1
1.97	0.39	69	RGB	0.15	0.08	1.0	0	u-grid(E)	0.7	250	1.3	4	1
1.97	0.39	69	RGB	0.3	0.15	1.0	0	u-grid(E)	0.7	250	1.4	10	1
1.97	0.39	69	RGB	0.6	0.3	1.0	0	u-grid(E)	0.7	250	2.4	16	1
1.97	0.39	69	RGB	0.9	0.5	1.0	0	u-grid(E)	0.7	250	4.6	11	1
0.88	0.39	85	RGB	0.1	0.11	1.0	0	u-grid(E)	1.7	1000	4.2	–	2
0.88	0.39	85	RGB	0.15	0.17	1.0	0	u-grid(E)	1.7	1000	4.7	–	2
0.88	0.39	85	RGB	0.3	0.34	1.0	0	u-grid(E)	1.7	1000	9	–	2
0.88	0.39	85	RGB	0.6	0.68	1.0	0	u-grid(E)	1.7	1000	16	–	2
0.88	0.39	85	RGB	0.9	1.02	1.0	0	u-grid(E)	1.7	1000	22	–	2
0.88	0.39	83	RGB	0.1	0.11	1.0	0	SPH(S)	500k	1050	5.7	2	2
0.88	0.39	83	RGB	0.15	0.17	1.0	0	SPH(S)	500k	950	7.8	6	2
0.88	0.39	83	RGB	0.3	0.34	1.0	0	SPH(S)	500k	750	10	8	2
0.88	0.39	83	RGB	0.6	0.68	1.0	0	SPH(S)	500k	950	18	10	2
0.88	0.39	83	RGB	0.9	1.02	1.0	0	SPH(S)	500k	600	25	10	2

¹Stellar spin frequency as a function of orbital frequency. None of these simulations used spinning stars.

²u-grid: uniform, static grid; SPH: smooth particle hydrodynamics. E: ENZO; S: SNSPH.

³Rounded to 2 significant figures, calculated at the end of the simulation.

⁴Calculated by including thermal energy.

⁵1: This work. 2: P12.

TABLE 5.1: Initial conditions and final outcomes of the 512³ simulations carried out in this chapter, compared to the simulations of P12.




FIGURE 5.2: Initial configurations for the binary system used in this chapter ($M_1 = 1.97 M_\odot$) at different resolutions: 128^3 (top left), 256^3 (top right) and 512^3 (bottom).




FIGURE 5.3: Evolution of the orbital separation between the point-mass particles over the simulation time. Different colours represent different mass companions (see legend). The discontinuous look of these curves is due to the relatively low output data dump rate.

5.2 Evolution of the binary separations

Let us consider first the evolution of the orbital separation for our simulations (Figure 5.3). More massive companions sink deeper. Also, the initial slope changes substantially for different companion mass, with the least massive companion (Figure 5.3, black line) showing a slower decrease in the separation. The decrease steepens for more massive companion masses. The behaviour is exactly the same noticed in P12 (their Figure 4). Clearly the far greater binding energy of the more massive stars contributes to this behaviour.

The physics of the rapid in-spiral is dictated by the strength of the gravitational drag force (F_{drag}) exerted by the primary's envelope on the companion, which, for a subsonic regime is given by (Iben & Livio 1993):

$$F_{\text{drag}} \propto \frac{M_2^2 \rho v_{\text{rel}}^2}{(v_{\text{rel}}^2 + c_s^2)^2}, \quad (5.1)$$

while for the supersonic regime is given by (Hoyle & Lyttleton 1939):

$$F_{\text{drag}} \propto M_2^2 \rho v_{\text{rel}}^2. \quad (5.2)$$

Here M_2 is the mass of the companion, ρ is the gas density around the companion, v_{rel} is the relative velocity between the companion and the gas and c_s is the sound speed of the gas around the companion. The trend between speed of in-spiral and companion mass is expected in the initial phases of the in-spiral, since the gravitational drag depends, among other quantities, on the mass of the companion and relative velocity, both of which are larger for more massive companions. Examining the separation in Figure 5.3 later on, at approximately 150 days we see a regime change: the less massive companions sink deeper in the envelope of the primary than the more massive ones. Also, the more massive companions have separations that plateau earlier. The more massive companions expel the outer layers of the envelope right away, because of the sudden, larger energy injection. Hence in the case of a massive companion, starting from exactly the same conditions as a lighter one, the outer envelope is spun-up more at the beginning of the interaction, reducing the efficiency of the gravitational drag more rapidly and slowing the in-spiral. This said it is not always obvious why the in-spiral stops by plotting the quantities that make up the factors in Equations 5.1 and 5.2. Carrying out a study on such factors becomes increasingly more difficult as the companion mass increases, because, due to stronger forces close to the object, the fluid situation around it is more complex. For example, Staff et al. (2016b) tested the effect of planets interacting with the envelope of giant stars. For such low mass companions a local fluid analysis was possible, because the conditions around those objects were much less turbulent. Another difficulty we face here arises from the initial setup used. The envelope has to adapt to the companion as the simulation starts, resulting in additional complexity around the in-spiraling object. We attempt such a study for our lowest mass companion. In Figure 5.4 we compare the velocity of the companion, the local average gas velocity projected on the direction of the companion velocity ($\langle v_{\text{gas},\parallel} \rangle$) and the local average gas velocity perpendicular to the direction of the companion velocity ($\langle v_{\text{gas},\perp} \rangle$) in the left panel, the companion's Mach number and the normalised average gas density in the companion's proximity ($\langle \rho \rangle / \langle \rho \rangle_{\text{max}}$) in the right panel, for the least massive companion ($0.1 M_{\odot}$). Average quantities are computed over a sphere of $7 R_{\odot}$ (10 cells) centred around the companion. For further details we refer the reader to Section 3.3.4. The relative velocity between the companion and $\langle v_{\text{gas},\parallel} \rangle$ does not seem to be playing a role in the lack of initial in-spiral. Also, one may expect $\langle v_{\text{gas},\parallel} \rangle$ to increase as the gas is spun-up, but this clearly does not happen in this case. Finally, the relative velocity between companion and gas around it on the gravitational drag are the magnitude




FIGURE 5.4: Simulation with a $0.1 M_{\odot}$ companion. Left panel: companion velocity (thicker blue line), local average gas velocity projected on the direction of the companion velocity ($\langle v_{gas,\parallel} \rangle$, thick green line) and local average gas velocity perpendicular to the direction of the companion velocity ($\langle v_{gas,\perp} \rangle$, thin cyan line). The three lines are smoothed with a Savitzky-Golay filter, using 17 coefficients and 7th order polynomials. Right panel: companion's Mach number (thick blue line) and normalised average gas density in the companion's proximity ($\langle \rho \rangle / \langle \rho \rangle_{max}$, thin green line). The dashed vertical line marks the point of maximum density.

of $\langle v_{gas,\perp} \rangle$, always larger than $\langle v_{gas,\parallel} \rangle$. This indicates that the envelope gas tends to move outward radially instead of following the companion.

The lack of initial in-spiral must be then down to the low density in tandem with the low companion mass. After the 50 days mark the density around the companion starts increasing, until it plateaus at $\simeq 150$ days. In this time span the increasing density is the main cause of the increasing drag force which brings the companion deeper into the primary's envelope. Ostriker (1999), in their simulation of the gravitational drag, explore the effect of the Mach number. They show that near Mach number = 1 the gravitational drag is at maximum, decreasing steeply for values lesser than unity and decreasing more gradually for values larger than unity. Staff et al. (2016b) ascribed the end of the in-spiral to a change of the regime from supersonic to subsonic. In Chapter 3 we saw how the $0.6 M_{\odot}$ companion remains subsonic in its in-spiral; the reason for the end of the in-spiral is therefore not the Mach number. Here we observe instead that the Mach number of the companion is below unity to begin with, but the journey rapidly becomes supersonic. After peaking at $\simeq 50$ days, the Mach number decreases on the same time-scale, reaching its minimum around the sonic point. After the minimum the companion's journey becomes supersonic again, showing constantly increasing Mach numbers until the end of the simulation. This behaviour is difficult to interpret in connection to its effect on the gravitational drag and one has to wonder what effect a

supersonic speed has on a hydrodynamic simulation like ours. The formation of a shock is treated with an artificial viscosity factor that smooths the shock front over more than one cell. One may therefore be suspicious that the artificial treatment of the shock condition may influence the result. As we highlight its importance, we leave this topic to future work. We further posit that at low masses gravitational drag is not the only mechanism that should be driving the in-spiral. Staff et al. (2016b), working with planetary mass companions, showed analytically that the hydrodynamical drag on those bodies should be similar to the gravitational drag simulated in the code. This could also be the case for the least massive companions studied by P12 and here, a comparison we leave to future work.

Finally, we also wonder whether the hydrodynamic code does not simulate hydrodynamical drag. The most massive companions travel through the envelope together with a certain amount of gas that is trapped permanently in their potential well. This gas gives the companion an effective width which will slow the object down hydrodynamically. This effect may not impact the simulation outcomes, since it would be larger for more massive companions, which are also those with larger gravitational drag. However it would be useful to quantify it.

In Chapter 3 we have discussed how the definition of final separation is different for different publications. Some of them use a criterion based on reaching a certain percentage of the maximum orbital decay while others simply take the value at the end of their simulations. The choice of what to use is completely arbitrary, but one needs to be consistent when comparing. For this set of simulations we decided to take the values at $\simeq 250$ days, when the separation vs time curve has approximately flattened for all the simulations. The final separations reached in this work are plotted in Figure 5.5. All the final separations obtained by using a more massive giant are in lower and in better agreement with the observed samples. This behaviour is the result of a more compact and/or more massive primary. Looking at the results of these simulations one could think that this is the simple solution to the CE problem, however, pre-empting the results of Section 5.3, a more massive primary results in less unbound gas.

As a final remark, we note that, similarly to what was observed by P12, there is a correlation between the final separation and the mass ratio. This is not an observed correlation, although one may wonder whether observed systems are too heterogeneous a group to reveal such a correlation.




FIGURE 5.5: Final orbital separation vs mass ratio $q = M_2/M_1$ for observed post-CE systems (Zorotovic et al. 2010 and De Marco et al. 2011, black dots) and for simulations (Sandquist et al. 1998, green circles; Ricker & Taam 2012, cyan triangle - note that here we report the separation of the simulations of Ricker & Taam 2012 which is lower than reposted in Ricker & Taam 2008 where the in-spiral had not come to an end; the 256^3 ENZO simulations of P12 are shown as yellow squares; Rasio & Livio 1996, magenta pentagon; PHANTOM simulations, blue hexagons, see Section 3.4; Ohlmann et al. 2016, grey diamond; Nandez et al. 2015, pink cross; the results of the ENZO simulation described in Chapter 3, yellow star; the results of this chapter, yellow rectangles).

5.3 Envelope ejection

To evaluate the unbinding of mass we plot the evolution of bound and unbound gas mass both inside and outside the box in Figure 5.6. We did not plot the total mass inside the simulation domain since it closely resembles the bound mass plot.

The trend of both bound and unbound mass inside the simulation domain are conditioned by the domain size itself. Mass starts leaving the domain very rapidly after the beginning of the simulations. This is clear from Figure 5.6 (left panels). The trends are a function of, as expected, the companion mass; lighter companions push the envelope out of the box slower and in smaller amounts. In all the simulations we notice an initial event of unbinding, but the mass unbound is




FIGURE 5.6: Bound mass inside the simulation domain (top left), unbound mass inside the simulation domain (bottom left), cumulative bound mass leaving the simulation domain (top right) and cumulative unbound mass leaving the simulation domain (bottom right). Different colours represent different mass companions (see legend).

entirely pushed out of the simulation domain very rapidly.

Of more interest is the right column of Figure 5.6, that shows the cumulative bound and unbound mass leaving the domain. We cannot store quantities that leave the domain during the simulations. Therefore the only way to estimate the bound and unbound mass leaving the domain is by interpolating between code data dumps as explained in detail in Section 3.3.2. The bound mass leaves the domain (Figure 5.6, top right panel) in a steady flow in all simulations, both during and after the rapid in-spiral phase. This is in accordance with the behaviour of the bound mass inside the simulation box discussed above (Figure 5.6, top left panel). Additionally, if for example we compare the yellow lines in Figure 5.6, top right and left panels, it is possible to see that the amount of bound mass inside the box decreases by $\simeq 0.4 M_{\odot}$ (there is almost no unbound mass in

the box at that time), a value that is the same as the sum of bound and unbound mass that left the domain, this shows the consistency of the interpolation method.

The total amount of mass unbound in the simulations has been calculated here by including the contribution of thermal energy; a more detailed discussion on the different contributions that can be considered is carried out in Section 3.3.2. We see from Figure 5.6 (see also Table 5.1) that in none of the simulations the entire envelope is unbound. In addition less massive companions unbind less envelope mass. The main unbinding event happens later for the lower companion masses and it is more gradual. The more massive companions are able to accrete and drag more envelope gas when the interaction starts, generating a wave of outgoing material with velocities larger than the escape speed, resulting in unbinding. For less massive companions the unbinding happens at later times, when the gravitational drag becomes more effective (as discussed in Section 5.2) and it is more gradual and less conspicuous.

This trend is not quite preserved by the $0.6 M_{\odot}$ companion, which is able to unbind 4% more mass than the more massive, $0.9 M_{\odot}$, companion. The reason is not immediately clear, but could be related to what is happening after the rapid in-spiral. From Figure 5.6 (bottom panels) it is in fact clear that unbinding happens after 100 – 150 days from the beginning of the simulation, and at this time the rapid in-spiral for the $0.6 M_{\odot}$ companion has already completed. While the initial unbinding of the envelope takes place with the first wave of material accelerated by the companion and it is mainly due to gas reaching escape velocity (Figure 5.7, left plot), the second one is mainly due to the heating of the gas, differently from what happens during the rapid in-spiral (Figure 5.7, right plot). These unbound regions coincide with zones where gas has a velocity larger than the local gas sound speed (Figure 5.8, top right panel). However, similarly high Mach number regions occur also for less (Figure 5.8, top left panel) and more (Figure 5.8, bottom panel) massive companions at late times and are not indicative of additional unbinding. There is no clear physical cause for this particular effect and it could be either a random occurrence or a numerical effect. A deeper understanding of what is taking place in the code is clearly necessary.

Overall, using a massive giant does not help increase the amount of mass unbound in the interaction with respect to a less massive one like the one used in P12, which achieves at its maximum an unbinding of $\simeq 10\%$ of the envelope.

5.4 Conclusions

In this chapter we performed a set of CE simulations using ENZO with an initial setup similar to the simulations of P12, except for the mass of the primary star’s envelope. The giant star we used is a $1.97 M_{\odot}$ star instead of a $0.88 M_{\odot}$ one, therefore it is more compact and denser for similar core masses. With such a setup we show that the entire set of companions, with masses 0.1, 0.15, 0.3, 0.6 and $0.9 M_{\odot}$, is able to achieve final orbital separations within the range of observed values (see Table 5.1). Similarly to P12, we obtain larger separations for increasing companion mass, something not readily observed. Finally, the final separations obtained by us and P12 for the lowest mass companions are much closer to each other than those obtained for higher companion masses. This suggests that at lower mass the in-spiral is more independent of the envelope mass and distribution than at higher companion masses. This has to be investigated in future work.

Although final separation is in agreement with observations, we are still unable to unbind the entire envelope mass. As expected, lower mass companions have less energy and angular momentum to transfer to the envelope and therefore unbind less mass than the heavier ones. The main unbinding event happens earlier, and in a shorter burst, at higher masses. This is due to the higher efficiency of the gravitational drag at early times. As the companion mass increases, the companion’s interaction with the envelope is less gradual, the envelope is lifted more rapidly and when the companion reaches the zones near the core the gas has a lower density. The percentages of unbound gas reach a maximum of 16% in the case of the $0.6 M_{\odot}$ companion, something that we are, for now, not able to explain. If we exclude this result the most mass unbound is 11% of the envelope mass, not dissimilar from the results of P12 and slightly lower than what we obtained in Chapter 3 and Chapter 4. We think that an additional physical mechanism not included in simulations is needed to match the observations. This could be the injection of recombination energy as gas recombines in the expanding and cooling envelope (Nandez et al. 2015).

Aside from code shortcomings it is not unrealistic to suppose that relatively more post-CE binaries derive from more massive stars, which have a larger binary fraction, something that we discuss further in Chapter 7.




FIGURE 5.7: Left columns: density slices perpendicular to the z axis in the orbital plane after, from top to bottom, 15, 29, 44, 58 days (left plot) and 234, 277, 292, 307 days (right plot) for the $0.6 M_{\odot}$ companion. The point-mass particles representing the core of the primary and the companion are shown as black dots, while the white regions represent the unbound gas. The size of the black dots is not representative of any property of the point masses and is chosen only to highlight them. Right columns: same as the left column, but excluding thermal energy from the computation of the bound/unbound mass elements.




FIGURE 5.8: Gas Mach number slices perpendicular to the z axis in the orbital plane at the same times of Figure 5.7, for the $0.1 M_{\odot}$ (top left), $0.6 M_{\odot}$ (top right) and the $0.9 M_{\odot}$ (bottom) companions. Time increases from left to right and from top to bottom. The point-mass particles representing the core of the primary and the companion are shown as black dots, their size is not representative of any property of the point masses and is chosen only to highlight them. The blue contours represent Mach number equal to unity.

6

Energy conservation in ENZO

During the work described in Chapter 3 we discovered an energy conservation issue with the ENZO code as applied to our problem. Energy non-conservation can create non-physical behaviours in the simulated systems, hence yielding incorrect results for the final configuration of the binary (final separation, unbound mass, etc.). Here we investigate the cause of this problem.

Due to this work, which resulted in a viable solution to the problem that was affecting other simulations that were being carried out within our group, I became co-author of the paper by Staff et al. (2016a).

6.1 The problem of energy conservation in grid codes

In all the codes, whatever their numerical method (e.g., grid or SPH), conservation of physical quantities is a fundamental requirement. If we consider grid codes (for a general idea on how they

work see Section 2.1.1), they are known to have limitations in conserving energy, and this is in part due to how they do the integration of the energy equation. To carry out a numerical integration it is in fact convenient to insert gravity as a source term. ENZO is no exception. The numerically-solved energy equation can be expressed, in the pure hydrodynamical form, as (Jiang et al. 2013):

$$\frac{\partial e}{\partial t} + \nabla \cdot [(e + P)\vec{v}] = \rho \vec{v} \cdot \nabla \Phi, \quad (6.1)$$

where e is total energy per unit volume, P is pressure, t is time, \vec{v} is the velocity vector, ρ is density and Φ is gravitational potential. This means that gravity is only taken into account on the right hand side of Equation 6.1 and the gravity solver is separated from the hydrodynamic solver. In other words, the left hand side of the equation is solved by the hydrodynamic solver, while the right hand side is solved by the Poisson solver. As a results of this technique, at each time-step the quantities on the left hand side and those on the right hand side are not guaranteed to be the same to round-off error.

This problem is exacerbated if gravity cannot be evaluated with enough precision, for example when resolution is low and steep, moving density gradients are not resolved by a sufficient number of cells. Let us take as an example Figure 6.1, where the problem is explained in one dimension for simplicity. The blue curve represents the analogue density resolved by two cells, such that we have only two values for the density ρ at each cell centre. Here we call ρ_1 and ρ_2 as the values at the cell centres, but they can assume any value of ρ_{analytic} (the blue curve in Figure 6.1) within the cell boundaries, depending on the interpolation method. ENZO evaluates the potential energy at each cell centre on the base of ρ_1 and ρ_2 , values that have a standard deviation that is proportional to the cell size and the steepness of the gradient. Under these circumstances, when ENZO steps forward by one time-step, if mass moves between the two cells, the potential energy change at those cell centres has an error associated with the standard deviation of the density value.

As a consequence, at each successive time-step, the value of the right hand side of the energy equation is much more approximate than for the left hand side, hence yielding an erroneous value for energy. This situation, in our simulations, systematically occurs in the central regions of the primary. Here the density gradient is very steep and moving from one cell to another because of both the expansion of the star and its orbital velocity. Additionally we are not able to reach a




FIGURE 6.1: One dimensional scheme of a steep density gradient not resolved in a sufficient number of cells.

resolution such that the gradient is split between a sufficient number of cells.

The issue involving energy conservation we encountered in ENZO revealed itself in our simulations in view of the specific conditions encountered in CEs.

6.2 Evaluating energy conservation

Conservation of energy and angular momentum can be difficult to evaluate in grid codes when simulating problems that have different spatial scales involved. In fact one has to find the right compromise between resolution and domain size to properly reproduce the physics of a problem within reasonable computational time. Ideally, we would like to model the CE ejecta, for which spatial scales of tens of AU are needed. At the same time we need to resolve the inner part of the in-spiral on spatial scales of less than a solar radius. The compromise is to simulate a few AU and let some mass leave the computational domain.

Unfortunately, when mass leaves the simulation domain both energy and angular momentum are irreparably lost. Additionally, our initial configuration includes an extremely hot ($\simeq 10^8$ K), low density ($\simeq 10^{-11}$ g cm $^{-3}$) medium outside the giant, intended to keep the star from expanding in the otherwise empty space (see Section 2.2). The problem with this method is that, even before

mass from the primary is lost from the grid, the external medium is pushed out of the box by the initially expanding giant. This has almost no effect on the total mass and angular momentum in the computational domain. However, this gas carries a large amount of thermal energy because of its high temperature. Therefore, as soon as a small amount of gas leaves the domain the total energy in the computational domain is greatly affected. Obviously this issue created a problem when assessing energy conservation and could have therefore have partly affected our conclusions despite our effort to account for natural energy losses out of the domain. Alternative ways to evaluate energy conservation in grid codes include the method we have adopted in Section 3.3.2. The work carried out here, however, does not attempt to estimate the amount of energy leaving the domain. Here we account for the energy leaving the domain only qualitatively, something that, as it turns out, was sufficient to identify the origin of the problem and suggest a solution.

6.2.1 Calculation of energies

Here we rapidly describe how the computation of the various components of the total energy showed in the plot in Section 6.4 has been carried out. The total energy is given by the formula $E_{\text{tot}} = E_{\text{therm,tot}} + E_{\text{kin,tot}} + E_{\text{pot,tot}}$ and its single components are computed as follows.

The contribution to thermal energy comes just from the gas and its value is given as an output by the ENZO code per unit mass, hence for each cell it is just multiplied by the gas mass in the same cell. Then the contributions of all the cells are summed up to get $E_{\text{therm,tot}}$. The total kinetic energy is given by the sum of the kinetic energy of the gas and of the particles. The contribution of the gas is given by

$$E_{\text{kin,gas}} = \frac{1}{2} \sum_{i=1}^N m_i v_i^2, \quad (6.2)$$

where the index i runs over all the cells, N is the total number of cells, m_i is the mass of the i -th cell and v_i is the absolute value of the velocity of the i -th cell. The contribution of the particles is instead given by

$$E_{\text{kin,part}} = \frac{1}{2} \sum_{k=1}^M m_k v_k^2, \quad (6.3)$$

where the index k runs over all the particles, M is the total number of particles (1 for the cases where there is a single star or 2 when there is a binary), m_k is the mass of the k -th particle and v_k

is the modulus of the velocity of the k -th particle. Then $E_{\text{kin,tot}} = E_{\text{kin,gas}} + E_{\text{kin,part}}$.

The smoothed potential, Φ , of the gas at each cell position is given as an output by the ENZO code and the total potential energy between gas and gas has been computed as:

$$E_{\text{pot,gas-gas}} = \frac{1}{2} \sum_{i=1}^N \Phi_i m_i, \quad (6.4)$$

where Φ_i is the gas potential at the position of the i -th cell. The $1/2$ factor is to prevent a double contribution by each pair of cells. The total potential energy between particles and gas has been computed as:

$$E_{\text{pot,part-gas}} = \sum_{k=1}^M \sum_{i=1}^N \frac{-G m_k m_i}{\sqrt{r_i^2 + \epsilon^2 \delta^2 \exp[-r_i^2 / \epsilon^2 \delta^2]}}, \quad (6.5)$$

where G is the gravitational constant, r_i is the distance between the i -th cell and the k -th particle, $\epsilon = 1.5$ (this is a numerical multiplying factor which combined with δ gives the smoothing length) and δ is the size of a cell. Finally, the total potential energy between particles has been computed only in the case where there are two particles in the grid and it is given by

$$E_{\text{pot,part-part}} = \frac{-G m_1 m_2}{\sqrt{r_{12}^2 + \epsilon^2 \delta^2 \exp[-r_{12}^2 / \epsilon^2 \delta^2]}}, \quad (6.6)$$

where m_1 is the mass of the first particle, m_2 is the mass of the second particle and r_{12} is the distance between the two. Note that in the Equation 6.5 and Equation 6.6 the smoothed potential recipe from Ruffert (1993) has been used to correctly compute the energies.

6.3 Estimating the extent of energy non-conservation

Before we started using the AMR version of ENZO as implemented by Passy & Bryan (2014), we ran simulations with a static, uniform grid. In fact, the study of Chapter 3 had been initially carried out in this way. The simulation used a domain size equal to the one used in Chapter 3 ($863 R_{\odot} = 4 \text{ AU}$), but we doubled the number of cells with respect to P12, from 128^3 to 256^3 , to maintain the primary star resolved with the same number of cells as P12. The smoothing length for the point-mass particles was 3 times the cells size instead of the value of 1.5 used in Chapter 3




FIGURE 6.2: Components of the energy inside the computational domain over the whole simulation time: total energy (stars blue), total kinetic energy (solid green), total potential energy (dashed green), total thermal energy (dotted green), gas kinetic energy (solid red), gas potential energy (dashed red), gas thermal (dotted red), particles kinetic (solid yellow), particles to particle potential (dashed yellow) and particle to gas potential (dashed cyan).

and, as we will see, this factor was a main contribution to the problem of energy conservation. The remaining parameters were exactly the same as those used in Chapter 3, to which we refer the reader for further details.

In Figure 6.2 we show the various components of the energy inside the simulation domain over time. There is a steep rise in total energy (blue star line) in the first year of the simulation, followed by steep decrease and flattening. The decrease and flattening after the first year of simulation can be understood by looking at Figure 6.3, showing the total gas mass inside the simulation domain over time. At one year, significant mass starts leaving the box, a time that approximately corresponds to the point when the total energy decreases. Clearly, when a lot of gas is pushed out a commensurable amount of thermal energy is lost. This trend is also confirmed by how the total energy curve mimics the behaviour of thermal energy (dotted red curve in Figure 6.2). The subsequent flattening can also be explained by the trend of the thermal energy, in fact as soon as the “hot vacuum” has entirely left the domain, at around 2 yr, the decrease is much smoother and it is balanced by the increases in both potential and kinetic energies.

Let us now consider the more enigmatic initial increase in total energy, from $\simeq 1.1 \times 10^{46}$ erg to $\simeq 1.6 \times 10^{46}$ erg ($\Delta E_{\text{tot}} \simeq 5 \times 10^{45}$ erg). This cannot be accounted for by the change in thermal




FIGURE 6.3: Gas mass inside the domain over the whole simulation time.

energy, which goes in the opposite direction. Here we assume that during this time span all the thermal energy lost is due to “hot vacuum” mass leaving the box and only a negligible part has been transformed into kinetic energy. With this assumption all the thermal energy lost needs to be added to the energy non-conservation budget, making the total amount of energy change in the first 0.86 yr of the simulation 1.63×10^{46} erg. This is a substantial fraction of the total energy at the beginning of the simulation. It is also a substantial fraction of the potential energy in the system.

6.4 Further tests

We will now discuss in detail the set of test simulations we carried out to understand what was causing the energy non-conservation issue. Below we refer to each test using the reference label in the first column of Table 6.1. The tests are organised as follows: T01, T02 and T03 are tests done without a star, from T04 to T12 are tests done with only a giant star in the computational domain and from T13 to T16 are tests done by running CE simulations.

As a measure of the energy conservation we use the change in total energy, $\Delta|E|$, divided by the initial value of the total energy, $|E_{\text{initial}}|$. We acknowledge that the fractional change will be driven by how close to zero the system’s total energy is initially. As a result the value $\Delta|E|/|E_{\text{initial}}|$ is not necessarily an indication of how acceptable a given level of non-conservation might be.

TABLE 6.1: Test simulations list.

Reference	Resolution	Box size (cm)	System setup	Boundary conditions	Solver	Max $\Delta E / E_{\text{initial}} $	Cycles	Primary t_{dyn} (days)	Run time
T01	256	1.0	Sod Shock Tube	Periodic	Zeus	$\simeq 0.008$	56218	N.A.	1 s
T02	1024	1.0	Sod Shock Tube	Periodic	Zeus	$\simeq 0.007$	56218	N.A.	1 s
T03	256	1.0	Sod Shock Tube	Periodic	PPM	$\simeq 0.007$	56218	N.A.	1 s
T04	384^3	4.4×10^{14}	Single star moving ^a	Vacuum	Zeus	$\simeq 0.20$	38007	154	10 yr
T05	384^3	4.4×10^{14}	Single star static ^a	Vacuum	Zeus	$\simeq 0.035$	37670	154	10 yr
T06	256^3	4.4×10^{14}	Single star moving, ^a high density background	Vacuum	Zeus	$\simeq 0.14^c$	3141	154	10 yr
T07	256^3	4.4×10^{14}	Single star moving, ^a no central particle	Vacuum	Zeus	$\simeq 90$	2619	154	4.5 yr
T08	256^3	4.4×10^{14}	Single star static, ^a no central particle	Vacuum	Zeus	$\simeq 90$	2609	154	4.5 yr
T09	64^3	1.5×10^{13}	Single star static ^b	Periodic	PPM, dual	$\simeq 0.001$	67424	28	140 days
T10	64^3	1.5×10^{13}	Single star static ^b	Periodic	PPM, dual	$\simeq 0.07$	65600	24	140 days
T11	64^3	1.5×10^{13}	Single star static ^b	Periodic	PPM, dual	$\simeq 0.5$	65593	22	140 days
T12	64^3	1.5×10^{13}	Single star static ^b	Periodic	PPM, dual	$\simeq 1.6$	30627	6	140 days
T13	256^3	6.0×10^{13}	Binary 3 stellar radii ^d	Vacuum	Zeus	$\simeq 0.34^e$	56218^e	28	1 yr ^e
T14	256^3	6.0×10^{13}	Binary 3 stellar radii ^d	Periodic	Zeus	$\simeq 1.0$	55982	28	1 yr
T15	128^3	6.0×10^{13}	Binary 3 stellar radii ^d	Vacuum	Zeus	$\simeq 0.18^e$	49564^e	28	1 yr ^e
T16	64^3	6.0×10^{13}	Binary 3 stellar radii ^d	Vacuum	Zeus	$\simeq 0.34^e$	27016^e	28	1 yr ^e

^a $M \simeq 3.05 M_{\odot}$, $R \simeq 473 R_{\odot}$.

^b Simulations run with the same stellar models ($M_{\text{envelope}} \simeq 0.48 M_{\odot}$ and $R \simeq 100 R_{\odot}$) and different central particle mass: $0.0 M_{\odot}$ (T09), $0.2 M_{\odot}$ (T10), $0.3959 M_{\odot}$ (T11), $10.0 M_{\odot}$ (T12)

^c Mass is lost from the grid from the beginning of the simulation.

^d $M_{\text{primary}} \simeq 0.88 M_{\odot}$, $R_{\text{primary}} \simeq 100 R_{\odot}$, $M_{\text{companion}} = 0.6 R_{\odot}$.

^e Before the point when mass starts leaving the box.

T01, T02 and T03 solve the 1D Sod Shock Tube test with periodic boundary conditions to avoid mass loss and were run to check the energy conservation without gravity. Given the simplicity of this test we used it also to check the integrity of the Zeus solver we are using for CE simulations against the more computationally expensive PPM and to test resolution effects with the Zeus solver. Results show that with both solvers energy is well conserved and that energy conservation is worse with decreasing resolution.

T04 and T05, have been run with the AGB giant used by Staff et al. (2016a). The purpose of these tests was to analyse if the non-conservation changed depending on whether the star is static on the grid or moving. The comparison shows that clearly, in the moving case energy conservation is worse.

T06 aimed to analyse the effect of a higher density “hot vacuum” outside the giant. We suspected that non-conservation was generated at the edge of the giant, where it makes contact with the hot vacuum, because of the huge difference in the physical quantities of the two gasses. Also, we had previously discovered that a moving giant conserved less energy (T04), therefore we tested the higher density vacuum using the same moving star of T04, but reduced the resolution for computational time reasons. We therefore acknowledge that by changing the resolution, T06 is not perfectly comparable with T04 and T05. In this case the non-conservation appeared to be still high. However, a non-negligible flux of mass starts leaving the box from the beginning of the simulation, because in this case the vacuum is denser. This eliminates out of the domain more kinetic and potential energy compared to T04 and T05, but much less thermal energy. As a result conservation is better than T04, despite the resolution being lower.

T07 and T08 have been run to investigate how much the presence of the central point-mass particle affected the non-conservation. Here we removed the central point-mass. This resulted in the star expanding and losing mass out of the box. Here we highlight again that in Tables 6.1 we give the quantity $\Delta|E|/|E_{initial}|$ in column 7. This value is not always representative of good or bad energy conservation, because it is very sensitive to the value of the initial total energy (i.e., easily becomes very large if the initial total energy is close to zero). It can be therefore used as a comparison only between simulations with a similar initial total energy. This applies to both T07 and T08, where the initial value of the total energy is very close to zero, due to the absence of the point mass. As a result we cannot easily compare T07 and T08 with T04 and T05. We

can, however, compare T07 with T08. The comparison shows a similar amount of energy non-conservation, a sign that in the absence of a point-mass the presence of static or moving gas has no repercussions on the energy conservation. The issue is therefore connected with the point-masses.

We further investigate the issue with the point-masses from the two previous tests using T09-T12. This set of four tests uses the same giant star as in in Chapters 3 and 4. In this case we proceeded by trying different masses for the particle in the centre of the giant, to understand how much the issue depended on its actual mass. Therefore we also used an overall lower resolution, to reduce computational times. The results show that by increasing the point-mass value, energy conservation worsens, something that we ascribe to the fact that more massive particles create a steeper potential which would require higher resolution to be properly represented. Since the point-masses we use in these tests do not have the correct value to set the star in hydrostatic equilibrium, the giant is unstable and prone to expand and push mass out of the grid. To avoid this issue and correctly evaluate energy conservation we used periodic boundary conditions in these simulations.

We also run some tests with a CE simulation set up. T13 is the simulation in which we unveiled the energy non-conservation issue described in Sections 6.2 and 6.3, also visualised in Figure 6.2. In T14 we run exactly the same simulation as T13, but with periodic boundary conditions, to pin down the exact magnitude of energy non-conservation. The periodic simulation shows that the actual value of energy non-conservation, when no gas leaves the domain, is actually higher. Finally, we run two resolution tests, T15 (128^3 cells) and T16 (64^3 cells), to be compared with T13 and to check the behaviour of the code at different resolutions with point-mass particles in the grid. The boundary conditions were “outflow”, like for T13, to keep conditions the same. The result of this is that ENZO conserves less energy at decreasing resolution.

From all the tests described above, we can extrapolate three trends compromising energy conservation in our CE simulations: energy conservation improves with decreasing the mass of the point-mass particle, energy conservation is better when a particle is static on the grid, and energy conservation improves with resolution. In the next section we describe in detail the origin of the problem and our method to improve conservation.

6.5 Solving the energy conservation issue

Only recently Jiang et al. (2013) described a numerical solver able to better conserve energy by not taking gravity as a source term. In lieu of an expensive code modification, we resorted to reducing the steepness of the potential gradients by increasing the particle potential smoothing length (Equations 6.5 and 6.6). The smoothing length is a numerical method used to approximate the gravitational potential on short distance scales, where it rapidly decreases to minus infinity. To avoid the singularity and related numerical approximations, the natural potential function is modified so that within the smoothing length the potential is flat. As a result gravity is smoothed within this threshold distance. There are various ways in which gravitational potential can be smoothed. In our simulations, for the point-mass particles, we use the prescription of Ruffert (1993), also used by P12.

Staff et al. (2016a) tested various resolutions and concluded that even using AMR to increase the resolution in high density regions, we are not able to match the steep density gradient around the core with enough cells and be able to simulate the interaction in an affordable time and with tolerable energy conservation. By doubling the smoothing length, however, we reduced non-conservation. To double the smoothing length we modified the value of ϵ in Equations 6.5 and 6.6. The ϵ used by P12 and also in the simulation that displayed poor energy conservation (T13) was 1.5, therefore yielding a smoothing length $\delta\epsilon$ of 1.5 times the smallest cell size.

Increasing the smoothing length decreases the strength of gravitational interactions, potentially creating less accurate results. Staff et al. (2016a) tested the interplay between energy conservation and smoothing length with the aim of identifying the minimum point-mass particle smoothing length one can use and still keep energy conservation within reasonable limits and at the same time do not approximate too much the gravitational interaction. The result of their tests is that an optimal value for ϵ is 3 (a smoothing length of three times the smallest cell size). We have placed the entire paper in appendix.

By increasing the smoothing length from 1.5 to 3 times the length of the smallest cell in the domain we are able to conserve energy retaining a reasonable gravitational interaction. For example in the simulation used for the work described in Chapter 3, which is the same as T13, we are able to conserve energy to the 4% level over 19861 coarse grid cycles, corresponding to

260 days \simeq 0.71 yr. This estimate has been done in the initial portion of the simulation when the mass of external hot medium leaving the domain is negligible. The only difference between T13 and the simulation in Chapter 3 is that in the latter we use ENZO with AMR, a coarse resolution of 128 cells per side and two levels of refinement, which grants us double the resolution of T13 in the densest regions.

7

Conclusions and future work

Here we summarise the CE problem, our approach, numerical issues we faced and the results obtained. Additionally we outline new directions for future research on the topic.

7.1 An unsolved physical problem

The scientific topic investigated in this thesis is the common envelope (CE) interaction (Paczynski 1976, Ivanova et al. 2013), that falls within the range of phenomena called close binary interactions: at a certain point during the life of a binary star there is the possibility that the two components of the system approach each other enough to exchange mass. If the rate of mass transferred from one object (primary) to the second (companion) is too large for the companion to accommodate it, the envelope of the primary distributes itself around the binary, creating a CE. In this situation the companion in-spirals rapidly towards the core of the primary, reducing their orbital

separation and ejecting the envelope around them.

The CE phase is a short event in astrophysical time-scales (~ 1 yr), and up to now it has been observed for sure only once by Tytenda et al. (2011). The CE interaction is of great importance, as it is one of the few known channels to decrease a binary's orbital separation ($\lesssim 5 R_{\odot}$) and, as such, the only way to explain compact evolved binaries and mergers. Observations in fact show that the universe is populated by many cases of such post-CE systems (De Marco et al. 2011, Zorotovic et al. 2011). Additionally, post-CE objects evolve on and eventually result in a variety of additional phenomena, such as compact white dwarf, neutron star and black hole mergers, that produce gravitational waves, as in the case observed by Abbott et al. (2016a). They are also the systems that give rise to type Ia supernovae and some gamma-ray bursts. As a result, understanding the physics of the CE is crucial for modern stellar astrophysics.

Despite the fact that the basic concepts of the interaction were clear already at the time when Paczynski (1976) published his work, up to today the CE problem has not been understood enough for the theory to make viable predictions. The interaction involves multiple physical processes and, its inherently 3D nature, places it beyond the ability of analytical studies. The main way to investigate the CE interaction is therefore with 3D numerical simulations. Various studies have approached the problem numerically: Rasio & Livio (1996), Sandquist et al. (1998), Sandquist et al. (2000), Ricker & Taam (2008), P12, Ricker & Taam (2012), Nandez et al. (2014), Nandez et al. (2015), Ohlmann et al. (2016). However, despite the use of different numerical techniques (grid-codes, SPH, moving mesh) and different initial conditions (e.g., component masses, primary radius, initial separation), we are still not able to make viable predictions, nor relate the simulations to the observed post-CE systems. Namely, except for one single case which we will discuss in Section 7.1.1 (Nandez et al. 2015), simulations cannot simultaneously obtain a small final orbital separation and unbind the entire gaseous envelope.

7.1.1 On envelope recombination energy simulations

Nandez et al. (2015) introduced for the first time gas recombination energy in CE simulations. This physical mechanism supplies an additional energy source that helps unbind the envelope. By including recombination energy, Nandez et al. (2015) were able to fully unbind the envelope,

something never achieved before. Without recombination energy they unbound 50% of their envelope, while with recombination energy they unbound almost 100%.

However, S. Ohlmann (private communication) applied the same technique to their simulation (Ohlmann et al. 2016) and increased the fraction of unbound envelope from 8% to 16%. Clearly recombination energy makes a difference, but does not necessarily solve the problem. This raises questions about the validity of the Nandez et al. (2015)’s work. In particular, the system they simulated (a primary with a mass of $1.5 M_{\odot}$ and a companion with a mass of $0.36 M_{\odot}$) is different from any previous work, precluding the possibility of a direct comparison. Their resolution is also very low (200,000 particles) and no resolution test is presented. Finally, their paper contains one figure and almost no detail, nor any explanation, nor analysis of their simulation.

The paper by Nandez et al. (2015) came out late in my Ph.D. and, the code modifications required to include recombination energy would have required more time than was available for this thesis. We leave this investigation to future work (see Section 7.4).

7.2 The aim of this project

The numerical efforts aimed to understand CE (see Section 7.1) lack in coherence. Numerical techniques differ, convergence tests are few and far between, and comparisons are almost non-existent. Conclusions as to the nature of specific mechanisms are made based on individual simulations carried out with only one parameter setup. Notable exceptions are the work of Sandquist et al. (1998) and P12. Necessary comparisons with observations are often missing. Our aim is therefore to investigate the CE interaction with a systematic approach, that consists of isolating single physical mechanisms that could affect the outcome of CE and carry out numerical simulations specifically tailored to highlight that effect. We have used as a main term of comparison the publication of P12, who developed one of the codes we use. We have used different numerical schemes, grid-codes in both uniform and AMR modes (ENZO; O’Shea et al. 2004, Bryan et al. 2014) and SPH codes (PHANTOM; Lodato & Price 2010, Price & Federrath 2010), to maintain a handle on different outcomes due to the different numerics and also to exploit the advantages of each.

7.3 Results

The main CE interaction results under scrutiny are the final separation of the binary and the amount of mass ejected from the envelope. It is these two simulation outcomes that are compared across simulations where no more than one aspect is changed at any one time.

7.3.1 CE simulations with a rotating giant

We first report the results of Chapter 4, because they stand alone and support and enrich the discussion of the results of Chapter 3. How can the rotation of the primary's envelope influence the outcome of the CE interaction? Rotation injects additional kinetic energy and angular momentum into the primary, resulting in an envelope that is less bound. However, with a spinning envelope the strength of the gravitational drag driving the in-spiral is reduced.

Sandquist et al. (1998), the only ones to carry out a similar comparison, deduced that the final separation is similar whether rotation is included or not (4.4 vs. 4.7 R_{\odot}), while the unbound mass in the rotating star is 20% less compared to the static one, from 33 to 26% of the total envelope mass.

We first carried out a preliminary study to understand the feasibility of importing a model of a giant evolved with rotation. However, with 1D stellar evolution codes we were not able to achieve rotation velocities similar to those observed in giants in binary systems. Additionally, even having succeeded in obtaining a suitably fast-spinning giant, it would have been dubious how to map such a 1D model into the 3D computational domain. Therefore we approached the problem by artificially spinning-up our primary in the computational domain starting from the same stellar model as P12, calculated as a non-rotating star. The rotation velocity applied to the primary star was equal to 95% of the co-rotation velocity. No substantial difference is obtained by adding rotation to the envelope of the star. The final separation is very similar in both the rotating and non-rotating cases ($\simeq 10 R_{\odot}$), similar to the results obtained by Sandquist et al. (1998). The amount of envelope ejected is again very close for the two cases: the non-rotating setup unbinds 14% of the total envelope mass, while the rotating one unbinds 15% of it. This is in less close agreement with the results of Sandquist et al. (1998). From analytical estimates it is clear that the supply of angular momentum given by rotation is substantial, but the amount of kinetic energy

added is very low compared to the other energy reservoirs of the star. Hence it is not surprising that we do not see a noticeable change in unbound mass, in accordance with our result.

Repeating the $0.6 M_{\odot}$ companion simulation of P12 using the new ENZO AMR solver and a similar resolution resulted in somewhat different results. A final orbital separation of $10 R_{\odot}$ instead of $16 R_{\odot}$ and $18 R_{\odot}$ in the ENZO and SNSPH simulations of P12, respectively, and an unbound mass equal 14% of the envelope mass, compared to a 10% of the SNSPH simulation of P12 (see Section 7.3.1) reminds us that numerical and resolution differences do affect the outcome. The work of P12 include 500,000 particles SPH simulations as well as 256^3 uniform, static grid simulations. We leave to future work an accurate investigation of the possible resolution effects.

7.3.2 CE simulations with a wide initial orbital separation

We have analysed the effect of a large initial orbital separation on the outcome of the CE interaction by placing the primary and the companion at the maximum orbital distance that would allow a tidal capture. This differed from most of the initial conditions used in past simulations, that set the companion on top of the primary's surface, already overfilling its Roche lobe and triggering CE from the beginning of the simulation. With our setup we aimed to inject the energy and angular momentum from the orbit into the primary, in a way that is more in line with what happens in nature. We were not able to properly reproduce the tidal phase of the interaction, because the artificial setup of the binary stimulates oscillations in the primary's envelope that accelerate the tidal in-spiral. The subsequent phases of the interaction happen on time-scales comparable with theoretical expectations. Both final separation and unbound mass are larger than for the comparison simulations of P12. The final separation is twice as large compared to the results of P12 for the same mass ratio. The unbound mass is instead 14% larger. Therefore, a larger orbital separation favours more unbinding. However, we know that this is not due to the primary being spun-up (see Section 7.3.1), but rather to the different primary envelope geometry at the time of Roche lobe contact. The envelope unbinding is mainly due to thermal heating at early stages of the interaction, and to acceleration above escape velocity later on.

We carried out these simulations with both a SPH and an AMR grid-code. The results are in line with each other and the above trends are preserved when changing numerical scheme.

7.3.3 CE simulations with more bound primaries

The effect of the mass of the envelope on the outcome of the CE interaction is the last CE initial condition we considered in this work. A more massive giant star with identical core mass will leave a post-CE remnant that is indistinguishable from a less massive one, and yet we can expect the effects of a more massive envelope on the interaction to be substantial.

Here we have used a primary star with double the mass of the giant used by P12 and we have carried out a set of five different simulations at increasing companion's mass (the same set of companions as P12), where the companions are placed on the surface of the primary, as was the case in P12. By using a more massive primary star we observed a shortening of the time-scale of rapid in-spiral. This effect is due to the increased envelope density, which applies a greater gravitational drag to the companion as it plunges in. The evolution of the separation versus time curves for increasing companion's mass is similar to that observed by P12: more massive companions in-spiral faster, but their final separations are larger compared to the final separations reached by lower mass companions.

This difference in the rapidity of the in-spiral process is the reason why lower mass companions reach smaller final separations: more massive companions interact more strongly with early in the simulation, the gas density rapidly decreases and weakens the interaction, leaving more massive companions at larger orbital separations. This effect reduced for less massive companions, maintaining a higher gravitational drag and needing a deeper incursion of the companion to lift the envelope sufficiently to halt the in-spiral.

The values of final separations obtained for the more massive primary are in the range of the observed values for all the companion masses. We note that it is possible that observed systems do indeed derive from more massive primaries. Statistically, more massive stars are in binaries more often (Duchêne & Kraus 2013). Additionally, for the post-CE binaries with CO white dwarfs, implying that the CE interaction took place on the AGB, are more likely to form more massive stars: lower mass stars, with a small AGB vs. RGB radius ratio preferentially interact on the RGB (Soker 1998).

On the other hand, the percentage of envelope unbound in the interaction is similar to P12. The amount of mass unbound increases with increasing companion mass, to a maximum of 11%

for the $0.9 M_{\odot}$ one. This implies that although the mass of the primary star helps reducing the final separation, an additional source of energy remains necessary to fully unbind the envelope. Recombination of the envelope gas, as proposed by Nandez et al. (2015), could be a possible or partial solution to the problem.

7.3.4 Energy conservation in common envelope simulations

We have also investigated the energy conservation behaviour of the ENZO code applied to CE simulations. Under the physical regime of CE simulations ENZO tends to poorly conserve total energy. The issue was unveiled after performing simulations for the project of Chapter 3. To understand the origin of the problem we ran a systematic set of test simulations, isolating various possible factors. We found the issue to be the ENZO solver, that handles the gravitational potential as a source term in the energy equation. In other words, the Poisson equation is solved separately from the hydrodynamics. Therefore, moving, high density gradients are not sufficiently resolved and this generates approximations in the computation of the gravitational potential, that are added to the total energy at every time-step. The problem was reduced by carefully tuning the smoothing length of the point-mass to the resolution. Steep density gradients in our kind of simulations always occur close to the primary's core or companion. Therefore, reducing the strength of the potential close to the particle by increasing the smoothing length attenuates the errors produced. On the other hand, reducing the smoothing length also affects the realism of the gravitational interaction between particles and gas and the tuning process has to be checked carefully. The fruits of this work were exploited by Staff et al. (2016a), who carried out a series of test simulations, obtaining an optimal value for the smoothing length of 3 times the length of the smallest cell in the computational domain. This smoothing length is double that used by P12. This and the recent availability of AMR in ENZO as applied to CE (Passy & Bryan 2014) led us to repeat the P12 simulations used as comparison in this work.

7.4 Future work

In this thesis we explored only a couple of the possible initial conditions and physical mechanisms that could affect the outcome of the CE interaction. Below we give a list of all the work we will undertake next.

The work of Nandez et al. (2015) contains one of the most interesting approaches to CE simulations taken recently. As we have explained extensively in Section 7.1.1, this approach holds promise, but it is at the moment not clear that their results do not depend on the low resolution, or on the specific parameters used. Therefore it is essential to test this physical process further and with different stellar models, initial conditions and numerical codes. We are in the process of integrating MESA's equation of state both in ENZO and PHANTOM to test the impact of recombination energy in the simulation of P12.

Another likely important physical process, which has never been included in CE simulations is magnetic fields. Post-CE binaries inside planetary nebulae must have recently gone through a CE event, because planetary nebulae are short lived objects (De Marco 2009). Some of these planetary nebulae have jets that can be dated to before the CE ejection, while others have jets that have been launched after the CE interaction. To investigate these mechanisms, Tocknell et al. (2014) used velocities and spatial distributions of the jets to measure the strength of the magnetic fields present at the time of their launch. This is possible since jets and magnetic fields are intimately connected, and the former can be used as a proxy to understand the latter. For jets launched before the CE in-spiral, Tocknell et al. (2014) showed that the magnetic fields were weak at that time, of the order of 1 G. These jets are probably fuelled by an accretion disk that formed at the time of Roche-lobe overflow. Objects whose jets were launched after the CE in-spiral, show instead that the magnetic fields at that time were of the order of 500 – 1000 G, opening the possibility that a CE magnetic dynamo took place, winding up and intensifying the indigenous giant stars fossil magnetic field. This had been already predicted by Reg s & Tout (1995) and Nordhaus et al. (2007). Such a strong magnetic field could have an important dynamical impact on the interaction itself.

We have just started adding seed magnetic fields in a CE simulation, either as a constant, low strength field or as a dipolar field and stabilising the star with it. Once we find a way to create the initial conditions we will reproduce the simulation of P12. Additionally, it is important

to determine how the influence of magnetic fields changes with parameters such as primary and companion mass, particularly moving to a more massive regime.

If jets are launched before the CE event, their action could be important also during the inspiral. The idea of jets inside the CE has been introduced by Soker (2016). It would be interesting to introduce artificial jets in a CE simulation, starting at the moment of Roche-lobe overflow and ejected at the location of the companion, with ejection energy and momentum guided by observations and based on previous works (Staff et al. 2016a, Soker 2016). This could tell us if and how these jets modify the CE interaction, shape the morphology of the ejecta or alter the unbinding of the envelope and the final orbital separation reached.

Studies of CE interactions with massive stars are sorely lacking and are becoming increasingly important: the recent detection of gravitational waves from a double black hole merger demonstrates the existence of these systems. Double black holes and neutron stars that merge must have gone through at least one CE event. Understanding the CE phase has therefore become a must in the era of multi-messenger astrophysics. Using less massive and energetic CE interactions is of great importance to understand the dynamics of the problem, with the ultimate goal, however, to extend these simulations to massive stars to understand the formation of close binary neutron stars and black holes.

The CE interaction is also thought to be a relevant evolutionary stage of more complex stellar systems. As an example Tauris & van den Heuvel (2014) ascribe the formation of a triple star, where one component is a millisecond pulsar, to the interaction of a giant star with *two* white dwarfs. The three objects in the observed triple system are evolved stars, where nuclear burning is not active anymore. The age and type of the three stars is an indication that all of them went through phases where they were much more expanded in the past. This suggests that they have interacted during their evolution. Tauris & van den Heuvel (2014) propose a scenario where all the three components interact to form what we observe today. However, the orbital configuration is quite peculiar, with the millisecond pulsar in a close orbit (1.63 days period) with one white dwarf and the second white dwarf orbiting the pair with a much longer (327 days) period. A complex, multi-step interaction is therefore required to reach this configuration. The process described by Tauris & van den Heuvel (2014) is composed of nine steps, involving various physical mechanisms with different energetics, time-scales and length-scales. The first of these steps is the most critical.

Here a star (that is today's the millisecond pulsar) evolves to the giant phase and enters a CE interaction with a close companion, presumably a regular hydrogen-burning star. This CE phase could be reproduced using a massive giant ($M_{\text{MS}} \gtrsim 10 M_{\odot}$, only such massive stars later become neutron stars) and a nearby companion as well as with a second companion in a wider orbit. In this configuration the ejecta of the CE will interact with the wider companion and likely cause a gravitational drag on it that could remove energy and angular momentum from the wide orbit of this object and bring it closer than the observations show. The challenge will be to determine under what conditions the wide orbit can be preserved. Understanding this mechanism can be generalised to a range of systems relevant to the formation of double degenerate binaries composed of white dwarfs, neutron stars or black holes.

While trying to reproduce interesting and exotic systems is certainly a driver of theoretical common envelope studies, one must not forget that the range of physical mechanisms involved in the creation of such systems is likely beyond today's simulations. There remains therefore a need to continue the exploration of parameter space and of the effects of individual mechanisms on the interaction, alongside code development work that can increase the ability of the simulation tools. Aside from recombination energy and the action of magnetic fields, we should not forget that the final outcome of a CE interaction is dictated by a phase immediately following the dynamical in-spiral, which acts on a longer, thermal time-scale. This is unlikely to be reproduced in hydrodynamic simulations any time soon, although 1D simulations and partial 3D hydrodynamic simulations will be able to inform our intuition. Between the dynamical and the thermal phases one can also expect some bound envelope gas to return as a fall-back disk, something that is observed around certain post-CE binaries (van Winckel et al. 2009). The action of such disks may also alter the longer term outcome of the CE interaction. The path ahead is therefore long and yet a certain haste is dictated by the advent of time-resolved surveys and the new availability of gravitational waves to observe merging and interacting systems.



Appendix A: co-authored publications

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, *ApJ*, 818, L22
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, *Physical Review Letters*, 116, 061102
- Baglin, A., Auvergne, M., Barge, P., et al. 2006, in *ESA Special Publication*, Vol. 1306, *The CoRoT Mission Pre-Launch Status - Stellar Seismology and Planet Finding*, ed. M. Fridlund, A. Baglin, J. Lochard, & L. Conroy, 33
- Beck, P. G., Montalbán, J., Kallinger, T., et al. 2012, *Nature*, 481, 55
- Benacquista, M. J. & Downing, J. M. B. 2013, *Living Reviews in Relativity*, 16 [arXiv:1110.4423]
- Bond, H. E. 2000, in *Astronomical Society of the Pacific Conference Series*, Vol. 199, *Asymmetrical Planetary Nebulae II: From Origins to Microstructures*, ed. J. H. Kastner, N. Soker, & S. Rappaport, 115
- Bondi, H. 1952, *MNRAS*, 112, 195
- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Science*, 327, 977
- Bryan, G. L., Norman, M. L., O'Shea, B. W., et al. 2014, *ApJ*, 211, 19
- Cantiello, M., Mankovich, C., Bildsten, L., Christensen-Dalsgaard, J., & Paxton, B. 2014, *ApJ*, 788, 93
- Carroll, B. W. & Ostlie, D. A. 2006, *An introduction to modern astrophysics and cosmology*

- Ceillier, T., Eggenberger, P., García, R. A., & Mathis, S. 2012, ArXiv e-prints [arXiv]1210.5880]
- Darwin, G. H. 1879, *The Observatory*, 3, 79
- De Marco, O. 2009, *Publications of the Astronomical Society of the Pacific*, 121, 316
- De Marco, O., Passy, J.-C., Moe, M., et al. 2011, *MNRAS*, 411, 2277
- De Marco, O. & Soker, N. 2011, *Publications of the Astronomical Society of the Pacific*, 123, 402
- Deheuvels, S., Garcia, R. A., Chaplin, W. J., et al. 2012, *ApJ*, 756, 19
- Duchêne, G. & Kraus, A. 2013, *ARA&A*, 51, 269
- Eggenberger, P., Miglio, A., Montalbán, J., et al. 2010, *A&A*, 509, A72
- Eggenberger, P., Montalbán, J., & Miglio, A. 2012, *A&A*, 544, L4
- Eggleton, P. P. 1983, *ApJ*, 268, 368
- Endal, A. S. & Sofia, S. 1976, *ApJ*, 210, 184
- Fryer, C. L., Rockefeller, G., & Warren, M. S. 2006, *ApJ*, 643, 292
- Gafton, E. & Rosswog, S. 2011, *MNRAS*, 418, 770
- García-Segura, G., Langer, N., Różyczka, M., & Franco, J. 1999, *ApJ*, 517, 767
- García-Segura, G., Villaver, E., Langer, N., Yoon, S.-C., & Manchado, A. 2014, *ApJ*, 783, 74
- Goupil, M. J., Mosser, B., Marques, J. P., et al. 2012, ArXiv e-prints [arXiv]1211.1546]
- Herwig, F. 2000, *A&A*, 360, 952
- Hoyle, F. & Lyttleton, R. A. 1939, *Proceedings of the Cambridge Philosophical Society*, 35, 405
- Iben, Jr., I. & Livio, M. 1993, *Publications of the Astronomical Society of the Pacific*, 105, 1373

- Ivanova, N. & Chaichenets, S. 2011, *ApJ*, 731, L36
- Ivanova, N., Justham, S., Chen, X., et al. 2013, *The Astronomy and Astrophysics Review*, 21, 59
- Ivanova, N., Justham, S., & Podsiadlowski, P. 2015, *MNRAS*, 447, 2181
- Jiang, Y.-F., Belyaev, M., Goodman, J., & Stone, J. M. 2013, *New Astronomy*, 19, 48
- Kippenhahn, R., Weigert, A., & Weiss, A. 2012, *Stellar Structure and Evolution*
- Koch, D. G., Borucki, W. J., Basri, G., et al. 2010, *ApJ*, 713, L79
- Lodato, G. & Price, D. J. 2010, *MNRAS*, 405, 1212
- MacLeod, M. & Ramirez-Ruiz, E. 2015, *ApJ*, 803, 41
- Madappatt, N., De Marco, O., & Villaver, E. 2016, ArXiv e-prints [[arXiv:1608.03041](#)]
- Marques, J. P., Goupil, M. J., Lebreton, Y., et al. 2012, ArXiv e-prints [[arXiv:1211.1271](#)]
- Massarotti, A., Latham, D. W., Stefanik, R. P., & Fogel, J. 2008, *AJ*, 135, 209
- Miszalski, B., Acker, A., Moffat, A. F. J., Parker, Q. A., & Udalski, A. 2009a, *A&A*, 496, 813
- Miszalski, B., Acker, A., Parker, Q. A., & Moffat, A. F. J. 2009b, *A&A*, 505, 249
- Nandez, J. L. A., Ivanova, N., & Lombardi, J. C. 2015, *MNRAS*, 450, L39
- Nandez, J. L. A., Ivanova, N., & Lombardi, Jr., J. C. 2014, *ApJ*, 786, 39
- Nelemans, G. 2010, *Astrophysics and Space Science*, 329, 25
- Nordhaus, J. & Blackman, E. G. 2006, *MNRAS*, 370, 2004
- Nordhaus, J., Blackman, E. G., & Frank, A. 2007, *MNRAS*, 376, 599
- Ohlmann, S. T., Röpke, F. K., Pakmor, R., & Springel, V. 2016, *ApJ*, 816, L9

- O'Shea, B. W., Bryan, G., Bordner, J., et al. 2004, ArXiv Astrophysics e-prints [arXiv:astro-ph/0403044]
- Ostriker, E. C. 1999, *ApJ*, 513, 252
- P12: Passy, J.-C., De Marco, O., Fryer, C. L., et al. 2012, *ApJ*, 744, 52
- Paczynski, B. 1976, in *IAU Symposium, Vol. 73, Structure and Evolution of Close Binary Systems*, ed. P. Eggleton, S. Mitton, & J. Whelan, 75
- Palacios, A. & Brun, A. S. 2007, in *IAU Symposium, Vol. 239, IAU Symposium*, ed. F. Kupka, I. Roxburgh, & K. L. Chan, 431–436
- Passy, J.-C. & Bryan, G. L. 2014, *ApJ*, 215, 8
- Passy, J.-C., Mac Low, M.-M., & De Marco, O. 2012, *ApJ*, 759, L30
- Paxton, B., Bildsten, L., Dotter, A., et al. 2010, *MESA: Modules for Experiments in Stellar Astrophysics*, Astrophysics Source Code Library
- Paxton, B., Bildsten, L., Dotter, A., et al. 2011, *ApJ*, 192, 3
- Paxton, B., Cantiello, M., Arras, P., et al. 2013, *ApJ*, 208, 4
- Paxton, B., Marchant, P., Schwab, J., et al. 2015, *ApJ*, 220, 15
- Podsiadlowski, P. 2001, in *Astronomical Society of the Pacific Conference Series, Vol. 229, Evolution of Binary and Multiple Star Systems*, ed. P. Podsiadlowski, S. Rappaport, A. R. King, F. D'Antona, & L. Burderi, 239
- Politano, M. & Weiler, K. P. 2007, *ApJ*, 665, 663
- Price, D. J. & Federrath, C. 2010, *MNRAS*, 406, 1659
- Rasio, F. A. & Livio, M. 1996, *ApJ*, 471, 366
- Regós, E. & Tout, C. A. 1995, *MNRAS*, 273, 146
- Ricker, P. M. & Taam, R. E. 2008, *ApJ*, 672, L41

- Ricker, P. M. & Taam, R. E. 2012, *ApJ*, 746, 74
- Ruffert, M. 1993, *A&A*, 280, 141
- Sandquist, E. L., Taam, R. E., & Burkert, A. 2000, *ApJ*, 533, 984
- Sandquist, E. L., Taam, R. E., Chen, X., Bodenheimer, P., & Burkert, A. 1998, *ApJ*, 500, 909
- Soker, N. 1997, *ApJ*, 112, 487
- Soker, N. 1998, *ApJ*, 496, 833
- Soker, N. 2016, *New Astronomy*, 47, 88
- Springel, V. 2010, *ARA&A*, 48, 391
- Staff, J. E., De Marco, O., Macdonald, D., et al. 2016a, *MNRAS*, 455, 3511
- Staff, J. E., De Marco, O., Wood, P., Galaviz, P., & Passy, J.-C. 2016b, *MNRAS*, 458, 832
- Stone, J. M. & Norman, M. L. 1992a, *ApJ*, 80, 753
- Stone, J. M. & Norman, M. L. 1992b, *ApJ*, 80, 791
- Taam, R. E. & Ricker, P. M. 2010, *New Astronomy Reviews*, 54, 65
- Taam, R. E. & Sandquist, E. L. 2000, *ARA&A*, 38, 113
- Tauris, T. M. & van den Heuvel, E. P. J. 2014, *ApJ*, 781, L13
- Tocknell, J., De Marco, O., & Wardle, M. 2014, *MNRAS*, 439, 2014
- Turk, M. J., Smith, B. D., Oishi, J. S., et al. 2011, *ApJ*, 192, 9
- Tylenda, R., Hajduk, M., Kamiński, T., et al. 2011, *A&A*, 528, A114
- van Winckel, H., Lloyd Evans, T., Briquet, M., et al. 2009, *A&A*, 505, 1221
- Webbink, R. F. 2008, in *Astrophysics and Space Science Library*, Vol. 352, *Astrophysics and Space Science Library*, ed. E. F. Milone, D. A. Leahy, & D. W. Hobill, 233

- Zahn, J.-P. 2008, in EAS Publications Series, Vol. 29, EAS Publications Series, ed. M.-J. Goupil & J.-P. Zahn, 67–90
- Zorotovic, M., Schreiber, M. R., Gänsicke, B. T., & Nebot Gómez-Morán, A. 2010, A&A, 520, A86
- Zorotovic, M., Schreiber, M. R., Gansicke, B. T., et al. 2011, A&A, 536, L3