



Hydrodynamic simulations of white dwarf-white dwarf mergers and the origin of R Coronae Borealis stars

Downloaded from: <https://research.chalmers.se>, 2024-12-22 17:41 UTC

Citation for the original published paper (version of record):

Shiber, S., De Marco, O., Motl, P. et al (2024). Hydrodynamic simulations of white dwarf-white dwarf mergers and the origin of R Coronae Borealis stars. *Monthly Notices of the Royal Astronomical Society*, 535(2): 1914-1943.
<http://dx.doi.org/10.1093/mnras/stae2343>

N.B. When citing this work, cite the original published paper.

Hydrodynamic simulations of white dwarf–white dwarf mergers and the origin of R Coronae Borealis stars

Sagiv Shiber¹,^{1*} Orsola De Marco,^{2,3} Patrick M. Motl,⁴ Bradley Munson,¹ Dominic C. Marcello,^{1,5} Juhan Frank,¹ Patrick Diehl,^{1,5,6} Geoffrey C. Clayton⁷,¹ Bennett N. Skinner⁸,^{7,8,9} Hartmut Kaiser,^{5,10} Gregor Daib,^{5,11} Dirk Pflüger¹¹ and Jan E. Staff^{12†}

¹Department of Physics and Astronomy, Louisiana State University, Baton Rouge, LA 70803, USA

²School of Mathematical and Physical Sciences, Macquarie University, Sydney, NSW 2109, Australia

³Astronomy, Astrophysics and Astrophotonics Research Centre, Macquarie University, Sydney, NSW 2109, Australia

⁴The School of Sciences, Indiana University Kokomo, Kokomo, IN 46902, USA

⁵LSU Center for Computation and Technology, Louisiana State University, Baton Rouge, LA 70803, USA

⁶Applied Computer Science (CCS-7), Los Alamos National Laboratory, Los Alamos, NM 87545, USA

⁷Lunar and Planetary Laboratory, The University of Arizona, Tucson, Tucson, AZ 85721, USA

⁸Department of Physics and Astronomy, McMaster University, Hamilton, ON L8S 4M1, Canada

⁹Origins Institute, McMaster University, Hamilton, ON L8S 4M1, Canada

¹⁰Division of Computer Science and Engineering, Louisiana State University, Baton Rouge, LA 70803, USA

¹¹IPVS, University of Stuttgart, Stuttgart, D-70174 Stuttgart, Germany

¹²Department of Space, Earth and Environment, Chalmers University of Technology, SE-41296 Gothenburg, Sweden

Accepted 2024 September 26. Received 2024 September 18; in original form 2023 September 29

ABSTRACT

We study the properties of double white dwarf (DWD) mergers by performing hydrodynamic simulations using the new and improved adaptive mesh refinement code OCTO-TIGER. We follow the orbital evolution of DWD systems of mass ratio $q = 0.7$ for tens of orbits until and after the merger to investigate them as a possible origin for R Coronae Borealis (RCB) type stars. We reproduce previous results, finding that during the merger, the helium WD donor star is tidally disrupted within 20–80 min since the beginning of the simulation onto the accretor carbon–oxygen WD, creating a high temperature shell around the accretor. We investigate the possible helium burning in this shell and the merged object’s general structure. Specifically, we are interested in the amount of oxygen-16 dredged-up from the accretor to the hot shell and the amount of oxygen-18 produced. This is critical as the discovery of very low oxygen-16 to oxygen-18 ratios in RCB stars pointed out the merger scenario as a favourable explanation for their origin. A small amount of hydrogen in the donor may help keep the oxygen-16 to oxygen-18 ratios within observational bounds, even if moderate dredge-up from the accretor occurs. In addition, we perform a resolution study to reconcile the difference found in the amount of oxygen-16 dredge-up between smoothed-particle hydrodynamics and grid-based simulations.

Key words: hydrodynamics – methods: numerical – stars: evolution – binaries: close – white dwarfs – stars: carbon.

1 INTRODUCTION

R Coronae Borealis (RCB) stars are low-mass, hydrogen-deficient, carbon-rich giants, primarily made of helium (Clayton 2012). They are almost indistinguishable from a second class of stars, known as the hydrogen deficient, carbon (HdC) stars, with the exception that RCB stars are known to exhibit irregular and dramatic light variability in the form of deep declines that can leave the star at minimum for years before recovery is observed (Tisserand et al. 2022; Crawford et al. 2023).

For a long time, two scenarios were considered for their formation, a final helium shell flash in a single post-AGB star, or a merger of two WDs (e.g. Iben & Tutukov 1984; Clayton 2012). The discovery of very low $^{16}\text{O}/^{18}\text{O}$ ratios (Clayton et al. 2007; Karambelkar et al. 2022), much less than the solar value ~ 500 , along with the likely average mass of $\sim 0.9 M_{\odot}$ for these stars (Han 1998; Karakas, Ruiter & Hampel 2015) effectively eliminated the final flash scenario in favour of the merger one. Merging WDs are known to create ^{18}O , which, under certain circumstances, can be brought to the surface, making the $^{16}\text{O}/^{18}\text{O}$ ratio as low as unity (e.g. Crawford et al. 2020; Munson et al. 2021).

To determine whether these observations could be understood in the context of a WD merger, Staff et al. (2012) simulated WD–WD mergers of different mass ratios, $q = M_2/M_1 = 0.5 - 1.0$ with the hydrodynamic code FLOW-ER (D’Souza et al. 2006; Motl et al. 2007).

* E-mail: sshiber1@lsu.edu

† We are saddened to report the passing of our esteemed colleague and friend, Jan Staff.

In that paper, it was found that the correct oxygen ratio could only be achieved under specific conditions, and that mergers with mass ratio, $q = 0.7$, are the most likely to produce this ratio. A similar enrichment of ^{18}O has been found by Longland et al. (2011), when they analysed a merger of a $0.8 M_{\odot}$ CO WD and a $0.4 M_{\odot}$ He WD ($q = 0.5$) that had been simulated using a smoothed-particle hydrodynamics (SPH) code. Later on, Staff et al. (2018) carried out several additional simulations, using, in addition to the unigrid technique used by Staff et al. (2012), also an early version of the AMR technique code OCTO-TIGER (Marcello et al. 2021b), and the SPH code (SNSPH; Fryer, Rockefeller & Warren 2006). In that paper, they addressed the question of how much ^{16}O is dredged up from the accretor, CO WD, out to the surface of the merged star. The problem with too much dredge-up is that no matter how much ^{18}O is fused during the merger, its abundance will be greatly diluted by the dredged-up material. To reduce the dredge-up, Staff et al. (2018) also considered accretors that are hybrid WDs, having a thick shell of $> 0.1 M_{\odot}$ of helium on top of the CO core. They concluded that while FLOW-ER and OCTO-TIGER agreed, SNSPH produced a smaller amount of dredge up. Whether the accretor was a hybrid WD or not, had little effect on the results. The conclusion was that if OCTO-TIGER results were to be believed, we might be unable to reproduce the low $^{16}\text{O}/^{18}\text{O}$ isotopic ratio observed.

Following the work by Staff et al. (2012), Menon et al. (2013, 2019) carried out a study using a 1D implicit code to evolve a post-merger star and determine the abundance patterns. They used information from the hydrodynamics simulation in the form of conditions such as temperature in four radial zones. By applying a very specific mixing prescription, they could reproduce several observed characteristics of RCB star, including the $^{16}\text{O}/^{18}\text{O}$ ratio.

Later, Crawford et al. (2020) and Munson et al. (2021) carried out very similar studies. The latter study was based on a later 3D OCTO-TIGER WD merger simulation with a mass ratio $q = 0.6$, which they carried out specifically for their research. Instead of using a 4-zone system, as done by Menon et al. (2013), they mapped the 3D simulation into the 1D implicit code, subject to a stabilization phase to bring the object into equilibrium. They observed that it is much harder to obtain the correct observed abundance values in this way.

In this paper, we will use our new and improved OCTO-TIGER (Marcello et al. 2021b) to calculate several simulations of the $q = 0.7$ WD merger, similarly to what was done by Staff et al. (2018). This mass ratio is not only appropriate for an RCB star, but by using the same mass ratio as done previously it allows a greater ability to compare simulations. We carry out seven simulations, five OCTO-TIGER simulations, and two additional FLOW-ER simulations for code comparison. These simulations bracket in resolution the simulation of Staff et al. (2018) and allow us to study the convergence properties of the simulations and the amount of ^{16}O that is dredged up from the accretor. A second aim of this paper is to determine the temperature in the ‘Shell of Fire’ (SoF; Staff et al. 2012), the region in which partial helium burning can take place. This directly informs the amount of ^{18}O that is generated in the merger.

Finally, it is worth emphasizing that the existence of short-period double WD is beyond doubt and confirmed by numerous observations. Two commonly accepted evolutionary scenarios for the origin of these systems have been extensively explored in the literature (Iben & Tutukov 1986; Lipunov & Postnov 1988; Tutukov & Yungelson 1994; Li et al. 2019). For additional references to both observations and theory see Chen et al. (2022). The most likely origin for the systems we consider in this pa-

per is through mass transfer from a red giant onto a CO WD (Li et al. 2019).

In Section 2 we discuss details of the simulations that have been carried out in the past with $q \sim 0.7$. In Section 3 we describe our simulations’ setup. Section 4 shows in detail the time evolution of our reference simulation and discusses the numerical properties of all of our simulations. We investigate in Section 5 the properties of the merged object such as its spin, the nuclear reactions, and the ^{16}O dredge-up, and finally, in Section 6, we conclude.

2 PREVIOUS RELATED SIMULATIONS

In this section, we discuss previous WD–WD merger simulations of mass ratio $q = 0.7$, intending to integrate their results with ours. All previous comparable efforts have been summarized in the first part of Table 1.

The simulations that appear in Staff et al. (2012, 2018), Motl et al. (2017), and Diehl et al. (2021b), are listed in the first part of the table. The OCTO-TIGER simulation from Staff et al. (2018) also appear in Montiel et al. (2015). Three of the simulations were carried out with the SPH code SNSPH; the rest were carried out using the finite-volume codes, FLOW-ER (four simulations) and a previous version of OCTO-TIGER (five simulations, listed as ‘Octo’ in the second column). The system’s initial properties, initial total mass (M_{T}), initial orbital separation (a_0), and initial orbital period (P_0) are also listed. The equation of state (EoS) used in each simulation is shown in column 4, where ‘ZTWD’, ‘poly’, and ‘ideal’ stands for zero-temperature white dwarf, polytropic, and ideal gas equation of state, respectively. When using a polytropic EoS or an ideal gas EoS, there is one degree of freedom in converting the simulation units (from code units) to c.g.s. Therefore, the simulation units (and as such the system properties) can be scaled in those simulations, e.g. the simulations of Motl et al. (2017) and Diehl et al. (2021b). Using a ZTWD EoS, on the other hand, introduces two additional physical constants and thus the code units in these simulations cannot be scaled arbitrarily, and their conversion to c.g.s units is fixed. In Table 1, we scaled the initial orbital separation (column 9) and initial orbital period (column 10) of Motl et al. (2017) and Diehl et al. (2021b) to match the initial system properties of the ZTWD simulations closely.

In all previous $q = 0.7$ simulations the donor star was initially forced into contact with its Roche lobe by a ‘driving’ mechanism, specifically, by an artificial removal of angular momentum from the system at a rate of 1 per cent per orbit. This initiates a mass transfer that can be well resolved by the code early in the simulation time. Usually, driving was applied for several initial orbits. We list the exact driving duration for each simulation in column T_{drv} in Table 1. A longer driving time causes the system to achieve a deeper contact, which can alter the remaining evolution of the simulation.

The resolution of past simulations spans a relatively wide range. Comparing resolution across different techniques is not straightforward. For FLOW-ER, a uniform (non-AMR) cylindrical grid, we list the number of cells across each dimension, $N_R \times N_z \times N_{\theta}$, and the total number of cells, where R , z , and θ , denote the cylindrical radial, vertical, and azimuthal axes, respectively. For OCTO-TIGER, an AMR grid, we can list only the initial number of cells. This is a minimum value, as the number of cells increases as the simulation evolves in time. We also list in parentheses the smallest cell size in solar radii for OCTO-TIGER.

The AMR grid extends several times beyond the dimensions of the non-AMR grid and not all regions are equally resolved. A better comparison between the non-AMR and AMR codes is the number of

Table 1. Overview of all considered Q0.7 simulations. L12 is the simulation with 12 levels of refinement; the suffix ND means non-driven, and the suffix I means that the simulation was conducted in the inertial instead of the rotating frame; Simulations prefixed with FL are carried out with FLOW-ER.

Source	Code	M_{tot} (M_{\odot})	EoS	T_{drv} (P_0)	T_{merge} (P_0)	Resolution; Δx_{min} (R_{\odot}) (cells/particles) ^b	donor res (radial cells)	a_0 (10^9 cm)	P_0 (sec)	Hybrid
Staff 2012	Flow-er	0.9	ZTWD	2.0	10.2	226x146x256 = 8.4M	150	?	?	No
Staff 2018 ^a	Octo	0.89	ZTWD	4	≈ 22	?, 6.5×10^{-4}	?	3.48	?	No
Staff 2018	SNSPH	0.9	ZTWD	6?	?	20M particles	8.2M	?	?	No
Motl et al. 2017	Flow-er	0.9	poly	2.3	9.7	162x98x256 = 4M	70	3.414	114.0	No
Motl et al. 2017	Flow-er	0.9	poly	1.7	21.0	162x98x256 = 4M	58	3.414	114.0	No
Motl et al. 2017	SNSPH	0.9	poly	1.0	11.5	100k	85k	3.410	116.0	No
Motl et al. 2017	Flow-er	0.9	ideal	2.3	12.4	162x98x256 = 4M	70	3.414	114.0	No
Motl et al. 2017	SNSPH	0.9	ideal	1.0	10.0	100k	85k	3.410	116.0	No
Diehl et al. 2021b	Octo	0.85	ideal	2	6.7	0.5M; 1.5×10^{-3}	27	3.27	111.0	Yes
Diehl et al. 2021b	Octo	0.85	ideal	2	13.4	2.2M; 7.4×10^{-3}	54	3.27	111.0	Yes
Diehl et al. 2021b	Octo	0.85	ideal	2	15.5	3.1M; 3.7×10^{-4}	108	3.27	111.0	Yes
Diehl et al. 2021b	Octo	0.85	ideal	2	16.8	11.4M; 1.85×10^{-4}	215	3.27	111.0	Yes
This work										
L10ND	Octo (Bd)	0.9	ZTWD	0	36.8	1.7M; 1.2×10^{-3}	36	3.453	116.0	Yes
L11	Octo (Bd)	0.9	ZTWD	1.3	25.0	2.5M; 5.73×10^{-4}	73	3.413	114.0	Yes
L12	Octo (Bd)	0.9	ZTWD	1.3	38.7	5.3M; 2.86×10^{-4}	145	3.403	113.6	Yes
L12I	Octo (Bd)	0.9	ZTWD	0	11.9	5.3M; 2.86×10^{-4}	145	3.403	113.6	Yes
L13	Octo (Bd)	0.9	ZTWD	2.3	9.9	20.1M; 1.43×10^{-4}	290	3.401	113.4	Yes
FL-1	Flow-er	0.9	ZTWD	1	37	322x258x256 = 21.3M	130	3.413	114.0	No
FL-2	Flow-er	0.9	ZTWD	2	16.9	322x258x256 = 21.3M	130	3.413	114.0	No

^a See also Montiel 2015. ^b For the simulations done with OCTO-TIGER, in the Resolution column, we list the *initial* number of cells in the adaptive-mesh grid as well as the minimal cell width (in solar radii). ^c We have run, additionally, two non-driven simulations, L11ND and L12ND, not listed here, as they resulted in the accretor's expansion due to numerical instabilities (see Appendix D for more details). A question mark represents data that cannot be retrieved from available information.

cells across the donor's diameter, which we specify in the 'donor res' column of Table 1. For SNSPH, a SPH code, we list the total number of particles in the simulation (in the Resolution column), and the number of particles consisting of the donor only (in the 'donor res' column). The high number of particles in the donor relative to the total number of particles is meant to better resolve the mass transfer in those simulations.

In column 11 we specify whether the accretor is a hybrid CO-He WD rather than a CO-only WD. As mentioned, this can affect the amount of ^{16}O in the burning regions and consequently the ratio of ^{16}O to ^{18}O produced by the burning. We note that in Diehl et al. (2021b) the mass of the helium shell was $\sim 0.045 M_{\odot}$, while in the OCTO-TIGER simulations that we present in this paper (see Section 3.1) the helium shell is $0.13 M_{\odot}$. Lastly, we list the merger time, T_{merge} , which is the time from the beginning of the simulation until a merger occurred in the units of initial orbital periods.

From this table, it is notable that both the amount of driving and the simulation resolution affect the merging time. Specifically, the 4M-cell FLOW-ER run with polytropic EoS that has been driven for 2.3 orbits merged much faster compared to the other FLOW-ER run with the same EoS and the same resolution that has been driven only for 1.3 orbits. Alternatively, the simulations of Diehl et al. (2021b) were driven by the same amount. However, the merging time increases with finer resolution. This can be explained by the fact that longer driving pushes the system into deeper contact and higher initial mass transfer. As a result, the system evolves faster and eventually merges earlier. This is also true for higher resolution runs. The finer resolution promotes a lower mass transfer rate, and the evolution is slower. It is possible that at some fine resolution, the mass transfer rate could not be decreased below a certain value, which would represent the actual physical mass transfer rate, and the merging time will converge concerning resolution.

Besides affecting the merging time, resolution can have an effect on the temperature and on the different compositions of different atmospheric layers of the merger, in particular when considering a hybrid WD accretor. The temperature is critical for helium burning and the production of oxygen-18, while the composition is important for the mixing of oxygen-16 with oxygen-18. Examining the effects of resolution, together with extending to higher resolution, are some of the main goals of this paper.

3 SIMULATION SETUPS

In this section, we give details of the setups for all our new simulations, starting with those carried out with OCTO-TIGER and finishing with those carried out with FLOW-ER.

3.1 The OCTO-TIGER simulation setups

We carried out five simulations using the benchmarked version of OCTO-TIGER (Marcello et al. 2021b, labelled 'Bd' in Table 1) with four different resolutions corresponding to 10, 11, 12, and, 13 levels of refinement (L10 to L13 in Table 1). Simulation L13 has the highest resolution used for an OCTO-TIGER simulation, with almost twice the number of cells across the donor star than in the simulation by Staff et al. (2012). This simulation also has almost twice the initial number of cells than in the high resolution simulation of Diehl et al. (2021b). The initial number of cells and the smallest cell size in each run are listed in Table 1 (seventh column), as well as the number of cells across the donor WD (eighth column). The simulation domain extends up to ~ 200 times the initial separation, which allows us to follow the outflowing gas from the system up to a significant distance and for a longer time.

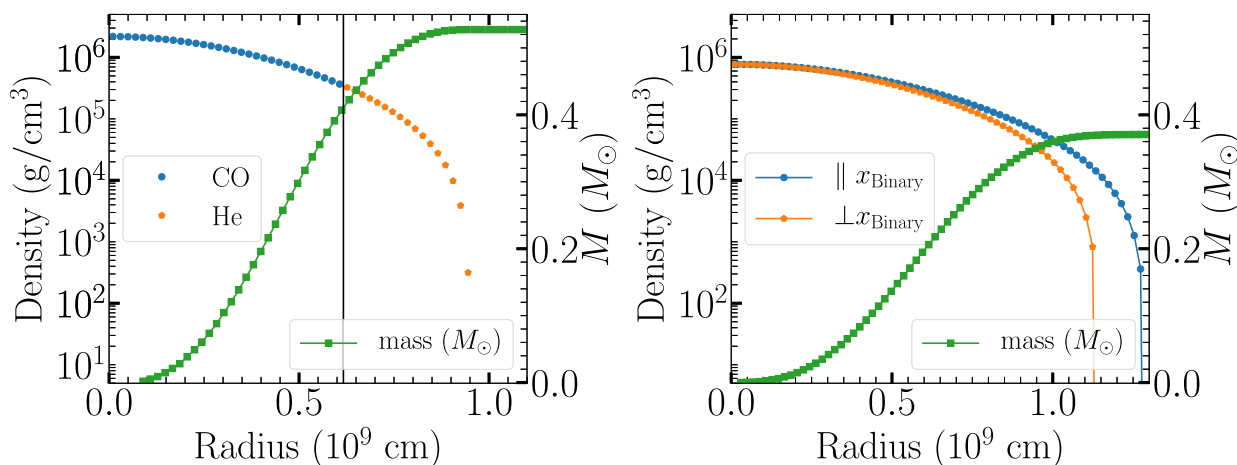


Figure 1. Initial density and mass profiles of the binary in simulation L12. Left: the hybrid accretor star, which is composed of a $0.4 M_{\odot}$ inner carbon–oxygen (CO) core and a $0.13 M_{\odot}$ outer helium (He) shell. Right: the donor star structure, where x_{binary} is the axis connecting the stars’ centres of mass.

Each simulation starts off with a WD–WD binary with a total mass of $M_{\text{T}} = 0.9 M_{\odot}$. The primary (accretor star) and secondary (donor star) initial masses are $0.53 M_{\odot}$ and $0.37 M_{\odot}$, respectively. We choose this value for the total mass because of observational indications (e.g. Karakas et al. 2015), and a mass ratio of $q = M_{\text{donor}}/M_{\text{accretor}} = 0.7$ is chosen because in previous simulations this ratio yielded the most suitable SoF temperature for incomplete helium burning (see Staff et al. 2012). The binary’s initial conditions are calculated using a self-consistent field (SCF) technique as described in Marcello et al. (2021b), except that a different equation of state, the zero-temperature white dwarf equation of state, is being used (see more details in Section 3.1.1). The donor star initially fills its Roche lobe and both stellar spins are synchronized to the orbital frequency. This results in an initial orbital separation of $0.05 R_{\odot}$ and in an orbital period of ~ 2 min. In all of the simulations except L12I, the grid is set to constantly rotate at the same frequency as the initial binary frequency. We carry out simulation L12I in an inertial frame of reference to study the impact of the chosen frame of reference on the system’s evolution and on the post-merger object’s properties.

The accretor is a hybrid He/CO WD, consisting of a $0.4 M_{\odot}$, CO core and an outer layer of $0.13 M_{\odot}$ of helium, while the donor is a helium WD. The CO WD core and its He envelope are evolved as two distinct fluids (or species) by OCTO-TIGER and the donor star is evolved as a third specified, distinct fluid. Each species has a mean molecular weight. The accretor CO core fluid is initialized with a mean molecular weight corresponding to an equal mixture of ionized carbon and oxygen. The mean molecular weights of the CO WD atmosphere and of the helium WD are $\mu = 4/3$, corresponding to ionized helium. The average mean molecular weight in a cell is used then to calculate the temperature.

In Fig. 1 we plot the resultant density and mass profiles of the accretor (left panel) and donor (right panel) stars, after the SCF initialization and for simulation L12. The inner $0.4 M_{\odot}$ of the accretor star ($3/4$ of the accretor’s mass) is the accretor CO core and it is surrounded by a He shell. The thickness of the accretor’s helium shell depends on resolution and is 5, 10, 18, and 36 cells, for the simulations with 10, 11, 12, and 13 levels of refinement, respectively. This shell serves as a buffer region between the CO core and the SoF, and therefore can potentially abate the dredge-up of ^{16}O from the core. Such buffer regions have been considered in

several previous numerical studies (e.g. Staff et al. 2018), and we choose a relatively massive shell to maximize its effect. However, in running our simulations we observe that this helium shell is almost entirely mixed in with the CO core well before the binary merges and therefore has only a minor effect on the resulting dredge-up of ^{16}O . This is mainly due to a numerical limitation of our code that cannot maintain a steep abundance gradient over an extended amount of time.

Lastly, to initiate a higher mass transfer rate which ultimately results in a faster merger, we extract angular momentum from the system’s orbit in simulations L11, L12, and L13 during the first 1.3, 1.3, and 2.3 orbits, respectively (fifth column of Table 1). This is done to save computer time. We drove L13 longer because in this highest resolution simulation, the initial mass transfer would be the lowest without enough driving and the computed time-step is also very small. The extraction of angular momentum is done by adding adequate source terms in the momentum equations as explained by Marcello et al. (2021b, their equation 44) and at a rate of 1 per cent of the orbital angular momentum per (initial) orbital period.

3.1.1 The equation of state

The construction of the stars by the SCF method uses a zero temperature white dwarf equation of state (ZTWD EoS), rather than a polytropic equation of state with an index $n = 3/2$. The ZTWD pressure is obtained by:

$$P_{\text{ZTWD}} := A \left[x (2x^2 - 3) (x^2 + 1)^{\frac{1}{2}} + 3 \sinh^{-1} x \right], \quad (1)$$

where $x := (\rho/B)^{\frac{1}{3}}$, A and B are constants, and ρ is mass density. Note that like the polytropic equation of state, equation (1) is also barotropic, allowing it to be incorporated into an SCF solver in a manner similar to what was done by Even & Tohline (2009).

After initialization, we evolve the simulations with a combined EoS of both ZTWD EoS and ideal gas (Staff et al. 2018):

$$P = P_{\text{deg}} + (\gamma - 1) E_{\text{th}}, \quad (2)$$

where E_{th} is the thermal energy density, and $\gamma = 5/3$ is the adiabatic index. The only difference with the method of Staff et al. (2018) is that we compute the local molecular weight for each cell rather than assuming it to be 2 for the temperature calculation, which results in




Figure 2. Conservation of mass, energy, and angular momentum in our simulations for $t \geq T_{\text{drv}}$. Vertical dashed lines denote the merger time. Note that in the angular momentum plot, we do not take into account angular momentum losses due to gas that flows out of the grid. P_0 is the initial orbital period of the simulation (between 113 – 116 s, depending on the specific simulation) as listed in Table 1.

lower temperatures compared to Staff et al. (2018). A more detailed description of this EoS and the exact calculation of temperature appears in Appendix A.

3.2 The FLOW-ER simulation setups

In order to compare the simulations’ outcome with a different code, we use the Eulerian uniform and cylindrical grid code FLOW-ER (D’Souza et al. 2006; Motl et al. 2007). We have carried out two additional simulations FL-1 and FL-2 with FLOW-ER. Both are evolved in the rotating frame and have identical resolutions. Simulations FL-1 and FL-2 are driven, by the removal of angular momentum uniformly at a rate of 1 percent, for 1 and 2 (initial) orbital periods, respectively (Table 1).

The density from the OCTO-TIGER SCF model is interpolated onto a uniform cylindrical mesh using cubic interpolation. Parameters from the OCTO-TIGER SCF model such as the angular frequency of the initial data are used with the same value as in OCTO-TIGER. The interpolation causes some initial oscillation in the density not present in the OCTO-TIGER runs but those perturbations decay well before the merger.

The spatial resolution has 322 radial zones, 258 vertical zones, and 256 azimuthal zones. The total number of cells is ≈ 20 M, larger than the 4 M cells resolution of the FLOW-ER simulations of Motl et al. (2017). The grid extends in radius to 6.4×10^9 cm and the spacing length in the radial direction equals to the spacing length in the vertical direction.

4 ORBITAL EVOLUTION AND MERGER

In this section, we start by describing the behaviour of the OCTO-TIGER simulation carried out with 12 levels of refinement in the rotating frame (L12 in Table 1). We later compare it to the other simulations.

4.1 Binary evolution leading up to merger

As mentioned before, in order to save computational time, the initial phases of simulation L12 were driven by continuously extracting angular momentum for the first 1.3 orbits at a rate of one percent per orbit. Still, the binary took 38.7 initial orbital periods, P_0 , until it merged. We simulated an additional five P_0 past-merger to allow the merger enough time to relax and to become more axially symmetric.

In Fig. 2 we present, three quantities that should be conserved, considering conservation properties only after the driving phase.

The quantities ΔM , ΔE , and ΔJ_z are computed as in Marcello et al. (2021b), and they include mass and energy that flows out of the grid, but not angular momentum outflows. For (L12), mass and energy were conserved to better than one part in 10^{-11} and 10^{-10} , respectively. Any inaccuracies stem solely from using a ‘density floor’, which prevents vacuum conditions in the low-density regions outside the star. We show in these plots the conservation values of simulations of different resolutions for comparison. We will discuss numerical effects like resolution in Section 4.2. The angular momentum is conserved at the 0.5 per cent level for L12, where most of the non-conservation stems from numerical viscous torques during and after the merger (as described in Marcello et al. 2021b).

In Figs 3 and 4 we show density and temperature maps, respectively, of the equatorial (top row) and meridional (bottom row) planes at three times, representing the beginning of the evolution (left panel), the time just before the merger (central panel), and the time of the merger (right panel). We also overlay velocity vectors (in the rotating frame) for gas with densities greater than 1 g cm^{-3} and equipotential contours. On the equatorial temperature slices, we overlay instead the grid structure (each square represents a subgrid of 8 by 8 by 8 cells) while on the meridional slices, we overlay density contours.¹

Shortly after the simulation begins, a stream of gas from the donor (less massive, larger star) flows through the L1 Lagrange point and flows around the accretor mainly along the equator. As the simulation progresses, mass starts flowing through the two outer Lagrange points as well and eventually, the donor is tidally stretched and wrapped onto the accretor. The cold donor material is heated by shocks when it impacts the accretor’s surface, and the temperature reaches helium-burning levels in a shell around the accretor. This SoF and its properties together with the potential nuclear fusion within it will be discussed in Section 5.2. In Fig. 5 we plot the overall orbital evolution on the orbital plane. The centre of mass position of each star is plotted, the accretor’s as a thick solid line and the donor’s as a thin dotted line. Colours represent time in units of the initial orbital period, P_0 (blue is the earliest, and dark red is the latest). To identify the cells of each star, we use the technique described in Marcello et al. (2021b). We plot the evolution until just before the merger, when this technique no longer works because the two individual stars are no longer well-defined. We typically find that the stars become indistinguishable by this technique when the binary orbital frequency, $\Omega_{\text{orb}} = j_{z,\text{orb}}/a^2$, lags behind the Keplerian frequency, $\Omega_{\text{kep}} = \sqrt{G(M_1 + M_2)/a^3}$, and we set the time of the merger, T_{merge} ,

¹Movies of the simulations can be obtained via [this link](#).




Figure 3. Simulation L12. Density slices at the equatorial plane (top row) and at the meridional plane (bottom row) with contours of Φ_{eff} . Arrows show magnitude and direction of velocities with respect to the instantaneous rotating frame, for gas with density greater than 1 g cm^{-3} . The time (in initial orbital periods, $P_0 = 113.6 \text{ s}$) is shown on the lower left corner of each panel.

according to the earliest time this occurs. For simulation L12 we obtain $T_{\text{merge}} = 38.7P_0 = 73 \text{ min}$.

Additionally, we plot in Fig. 5 the centre of mass of the entire grid in orange, and the combined centre of mass of the two stars in blue. The grid centre of mass remains in place, which indicates again a good degree of linear momentum conservation. The combined centre of mass of the two stars slightly shifts left during the merger, implying slightly asymmetrical outflows through the outer Lagrange points.

In Fig. 6 we show the orbital separation, the orbital angular momentum, the angular momentum of the primary (star 1), and the secondary (star 2), the mass of each star, and the mass transfer rate (in units of the total mass per initial orbital period), all as a function of time. During the first 1.3 orbits both the separation and the orbital angular momentum decrease, as expected due to the driving, while the orbital frequency increases and the binary axis rotates counterclockwise in the rotating frame.

After this initial driving phase, mass transfer onto the more massive, accretor star, acts to increase the separation. At this point, the mass transfer rate is $10^{-4} M_{\text{T}}/P_0 \simeq 3 \times 10^{-3} M_{\odot} \text{ h}^{-1}$ (lower right panel). However, the accretor also drains angular momentum from the orbit, at a rate of $10^{-3} J_{\text{orb},0}/P_0$ (upper-right panel), which counteracts the effect of mass transfer and tends to shrink the orbit. Overall the orbital separation remains fairly constant (upper-left panel). The mass transfer rate continues to increase, while the orbital angular momentum becomes angular momentum of the accretor star, and eventually, the donor star is tidally disrupted onto the accretor. This picture of unstable mass transfer that increases with time while

the orbital angular momentum decreases and that leads to the tidal disruption of the donor star is consistent with what is described in Motl et al. (2017) and is termed there a generic ‘case A’ evolution for WD–WD binaries. This is to be distinguished from a ‘case B’ evolution where the mass transfer stabilizes or decays with time, which leads to a dynamically stable binary configuration. We further discuss quantitatively the mass transfer and verify it against an analytic expression in Section 4.2.1 and in Appendix B.

Just prior, but mostly during the merger, mass is flowing through the outer Lagrange points, L2 and L3 at an increasing rate. Because our diagnostics method accounts gas as belonging to a star only if it is inside the star’s Roche lobe, we can quantify the amount of mass and energy lost from L2 and L3, ΔM and ΔE , as $\Delta M = (M_1 + M_2)|_{t=0} - (M_1 + M_2)|_t$, and $\Delta E = (E_1 + E_2)|_{t=0} - (E_1 + E_2)|_t$, respectively. In Fig. 7, we plot these quantities as a function of time. We find that a few per cent of the total mass is flowing out of L2 and L3, and that the system’s total energy becomes a few per cent more negative close to the merger. Therefore the energy outflow through L2 and L3 is positive. Most of the outflowing mass remains bound and ends up being accreted by the merged object at late times. For simulation L12 we find $\sim 1 \times 10^{-3} M_{\odot}$ of unbound mass, although we stopped the simulation before the amount of unbound mass reached a steady value. The amount of unbound mass decreases with resolution, and therefore the maximal value of unbound mass we find is in our lowest resolution run. We conclude that only a small fraction of the total mass becomes unbound as a result of the merger ($\lesssim 5 \times 10^{-3} M_{\odot}$).




Figure 4. Simulation L12. Temperature slices at the equatorial plane (top row) and at the meridional plane (bottom row) with, overlaid, the grid structure (top row), and density contours (bottom row). The time (in initial orbital periods, $P_0 = 113.6$ s) is shown on the lower left corner of each panel.




Figure 5. The orbital evolution of the system in simulations L12, showing the centres of mass of the individual stars until shortly before the merger, as well as the centres of mass of the system and of the grid – expected to be stationary. The stars merge at $t = T_{\text{merge}} = 38.7 P_0 \approx 4400 \text{ s} \approx 73 \text{ min}$.

To understand better the energy distribution prior to and during the merger, we plot in Fig. 8 the four energy components for each star and the system: kinetic (E_k , orange), gravitational (E_g , green), degenerate internal (E_{deg} , magenta), and thermal internal (E_{th} , red). The total for each of these four energy components for the system is shown as solid lines of the colour assigned above. The quantities due to the accretor (denoted by 1) are shown as dashed lines, while the corresponding contributions by the donor (denoted by 2) are

shown as dash-dotted lines. The total energy contributions for each component and the system are shown as thick blue lines. Note that the degenerate internal energy is calculated directly from the density using the corresponding expression for the ZTWD, while the thermal internal energy is calculated by a procedure described in detail in Appendix A. All of these energy quantities are measured in the inertial frame. The material flowing out of the donor is heated when it impacts the accretor at a high speed. This translates to a decline in the donor’s degenerate energy and an increase in the accretor’s thermal and kinetic energies. Eventually, all the energy contributions of the donor vanish at the merger. Interestingly, the donor material accreted onto the accretor still possesses some degenerate energy, and the total degenerate energy of the system after the merger is slightly greater than the initial accretor’s degenerate energy.

Fig. 9 shows the ‘virial error’, defined as $\text{VE} = |2E_k + E_g + 3\Pi|/|E_g|$, where E_k , and E_g , are the kinetic and gravitational energy, respectively, while $\Pi = \int P dV$ is the total pressure integrated over volume. The closer this quantity is to zero, the more accurate the numerical representation of the system. We plot the virial error as a function of time for simulation L12, although all other runs show the same trend as well. Initially, both stars individually obey the virial relation only approximately, because they are not isolated systems. The error is greatest for the donor (dashed-dotted line) since we are ignoring the more massive companion. As the donor approaches disruption, the relative virial error increases. In contrast, the virial error for the system is $\sim 10^{-3}$ throughout the evolution, and even after the merger (solid line), suggesting that our numerical representation of the system is very precise.




Figure 6. The orbital separation, angular momenta, masses, and mass transfer rates for the L12 simulation (see Table 1), all as a function of time. The system was driven for 1.3 orbits at a rate of one per cent of its orbital angular momentum per initial orbit and eventually merged at $T_{\text{merge}} = 38.7 P_0$. Note that the mass transfer is measured in units of total mass $M_T = 0.9 M_\odot$ per initial orbit, $P_0 = 113.6$ s for this simulation.




Figure 7. Mass (left panel) and energy (right panel) that flows out of the outer Lagrange points L2 and L3 in simulation L12. The energy of the merged object decreases as a result of the L2 and L3 outflow.

4.2 Numerical, physical, and conservation properties of the simulations

In this section, we discuss the results in light of the numerical properties of the simulations. As a first step, we have analysed our results during the early mass transfer to conclude the nature of the mass transfer (Section 4.2.1). Consecutively, we have tested several numerical properties of our simulations such as the effects of resolution (Section 4.2.2), of carrying out the simulation in the rotating and inertial frames (Section 4.2.3), and ultimately the amount of driving during the early part of the interaction (Section 4.2.4). For another numerical test, where we have found remarkably similar results across the different codes, OCTO-TIGER and FLOW-ER, we refer the reader to Appendix C.

4.2.1 Instability of mass transfer

The total angular momentum of a simple binary model, in which the component stars have constant masses and moments of inertia, and are tidally locked to the orbit, assumed circular, is the sum of orbital and spin terms at the same angular velocity. If the orbital-spin frequency is given by the usual point-mass Keplerian approximation, then the total angular momentum depends only on the binary separation a , and has a minimum at a separation $a_{\text{min}}^2 = 3(I_1 + I_2)/\mu$, where $\mu = M_1 M_2 / (M_1 + M_2)$ is the two-body reduced mass (Rasio 1995). With the initial moments of inertia of the equilibrium model binary, we find $a_{\text{min}} \simeq 0.48 a_0$.

If angular momentum losses drive such a binary to the above minimum, no further reduction of total angular momentum and




Figure 8. The evolution of the kinetic (orange), gravitational (green), thermal (red), degeneracy (magenta), and total (thick blue) energies for accretor (dashed), donor (dot-dashed), and for the binary system (solid) according to simulation L12. Note how the contributions of the donor vanish, and those of the accretor become the system's at the merger.




Figure 9. The virial error (VE) for the accretor (dashed), the donor (dot-dashed), and the binary system (solid). At the merger the VE for the donor diverges as E_g vanishes, and the VE for the accretor is essentially that of the system.

separation is possible unless synchronism is broken, and one or both components spin slower than the orbital frequency. But then tides will attempt to synchronize the lagging spins, reducing further the orbital angular momentum, and we have a run-away process, known as the Darwin instability (Darwin 1879). Therefore we expect that in a system driven by the Darwin instability, both orbital and spin frequencies increase with time, while the spin frequencies remain lower than the orbital. In all of our simulations the unstable behaviour begins well before a_{min} is reached (see below).

The simple arguments given above require revision if mass transfer begins at a separation exceeding a_{min} , if the mass is lost from the system, or if the tidal distortions of the stars are significant (Lai, Rasio & Shapiro 1994). All of these effects come into play in our simulations. For example, mass-loss through the outer Lagrange points,




Figure 10. The angular velocity of the orbit (Ω_{orb}), the accretor (Ω_1), and the donor (Ω_2) in the level 12 simulation (L12 in Table 1) divided by the system's initial angular velocity $\Omega_0 = 2\pi/P_0 = 3.32 \text{ rad min}^{-1}$ as a function of time. The system was driven for 1.3 orbits at a rate of one per cent of its orbital angular momentum per initial orbit.

and direct impact accretion by the CO WD, are consequential angular momentum losses (hereafter CAML) (Webbink 1984; King & Kolb 1995; Schreiber, Zorotovic & Wijnen 2016), meaning orbital angular momentum losses that are consequences of mass transfer. CAML and the expansion of the degenerate donor due to mass-loss, drive the mass transfer instability and ultimately cause the tidal disruption and the merger of the system. Because the accretor is spun up by mass transfer, its spin frequency increases faster than the orbital frequency.

In Fig. 10 we display the angular velocities for the orbit and the spins of the binary components. We see in this figure that the orbital angular frequency does not change much while the donor angular frequency decreases somewhat and the accretor spin frequency increases rapidly. This behaviour is a characteristic of mass transfer




Figure 11. Comparison between different resolutions. Shown are the orbital separation, orbital angular momentum, mass ratio, mass transfer rate, accretor, and donor central densities, all as a function of time.

instability and implies that the Darwin instabilities does not play any role in the merging process. We further analyse the binary evolution of our highest resolution L13 run, showing the excellent agreement between our simulations and the analytical expression in Appendix B.

4.2.2 The effect of spatial resolution

As can be seen in Fig. 2, OCTO-TIGER conserves mass and energy very accurately to the level of machine precision, regardless of the resolution being used. The source of non-conservation stems from the low values (floor values) of density and energy artificially introduced in a few cells to prevent the occurrence of vacuum conditions. In contrast, a finer grid does improve the conservation of angular momentum. For L10ND and L11 the angular momentum deviates by less than 4 percent by the time of the merger; however, the strong shearing during the violent disruption tends to contribute to conservation errors. Simulations L12 and L13 display a better

conservation of angular momentum at the level of less than a per cent even five orbits after the merger.

In Fig. 11 we present a set of diagnostic quantities plotted as a function of time for L10ND, L11, L12 (our reference simulation), L12I (same as L12, but performed in the inertial frame, see a detailed comparison in Section 4.2.3), and L13. Times have been shifted by T_{merge} so that the diagnostic quantities are shown lined up at the same time before the merger.

The smallest cell size in simulation L10ND is $1.2 \times 10^{-3} R_{\odot}$, which lets us resolve the donor diameter into 36 cells. Each additional level of refinement roughly halves the smallest cell size, and doubles the number of cells across the donor, so in L13 the donor diameter is resolved into 290 radial cells. Simulations L10ND and L12I were not driven, while the other three were driven, L11 and L12 for 1.3 orbits and L13 for 2.3 orbits, as described in Section 3. L12 was driven similarly to L11 and clearly merges later, as expected of a higher resolution simulation. So as to reduce the computational times, we have driven L13 by 2.3 orbits, which resulted in a faster merger.

The evolution of the orbital angular momentum is qualitatively the same in all our simulations, slowly declining at a rate of $10^{-4} - 10^{-3} J_{\text{orb},0}/P_0$ through most of the evolution (from the moment the driving has stopped until ≈ 5 orbits before the merging time), and then rapidly dropping at rate of $> 10^{-2} J_{\text{orb},0}^0/P_0$ just before the merger as the donor is tidally disrupted. During the tidal disruption, the orbital angular momentum is transferred to the spin of the accretor as discussed in Section 4.1.

The mass ratio, initially 0.7 for all simulations, decreases drastically close to the merger. In fact, the accretor in all of our simulations gains mass throughout the evolution, decreasing the mass ratio monotonically. L11 (dashed-dotted orange line) and L12 (solid green line) that were driven at the same rate and the same duration are the most similar, especially as illustrated by the evolution of the mass ratio q and the mass transfer rate \dot{M}_1 . While we expect a similar contact depth in both runs, L11 overestimates the initial mass transfer rate because of the larger cell sizes. Simulation L12 starts at a lower rate and takes around 20 orbital periods to ‘catch up’ with L11, but the rates agree very closely after that. Because the mass transfer rate is a very noisy quantity we smooth it using a Savitzky-Golay filter (Savitzky & Golay 1964) with a temporal window size equal to half of an orbit. As expected, the initial mass transfer is primarily a function of resolution, and higher resolution simulations yield lower initial mass transfer rates. None the less, the mass transfer rate during the final several orbits before the merger is very similar across all simulations.

As mass transfer proceeds, the expectation is that the central density of the accretor should increase, while that of the donor should decrease. The behaviour of the primary central density shows some variation initially but conforms to expectations as the mass transfer increases. The secondary’s central density behaves in a more homogeneous way: we observe in all simulations a decline just before the merger.

Our conclusions from the resolution study are as follows. First, resolution is important for an accurate angular momentum conservation. In addition, although the merging time itself depends on the initial depth of contact, which is a function of both resolution and amount of driving, the general behaviour before the merger is similar regardless of resolution. The WD–WD system merged in all of the simulations, regardless of resolution, supporting that WD–WD systems with mass ratio of 0.7 are unstable once mass transfer begins. This agrees well with previous simulations and analytical expectations (e.g. Motl et al. 2017). Lastly, in the following section we will show that resolution plays a critical role in determining the properties of the shell of fire (Section 5.2) and the amount of O^{16} that is being dredged-up during the merger (Section 5.3).

4.2.3 Inertial versus rotating frame simulations

To illustrate the advantages of conducting simulations in the corotating frame, we ran two simulations at the same resolution: L12 on the rotating frame, driven for 1.3 orbits, and L12I on the inertial frame, undriven. Simulation L12I involves advection of large flows across cells just to represent the orbital motion of the binary components, while mass transfer and internal motions in the binary components would be given by relatively small additional fluxes. Therefore we expect L12I to be prone to cumulative errors and larger diffusion.

Referring to Fig. 11 and focusing on the differences between the thick green curves (L12) and the thin magenta curves (L12I), we note that L12I ran to merger much faster than L12, even without

driving. The orbital angular momentum decreased from the start, and mass transfer was initially larger and increased rapidly. The central density of the accretor decreased rapidly for a couple of orbits from the beginning due to stronger numerical diffusion, and only began increasing once the mass transfer reached levels comparable to those present in L12, at $\sim 8P_0$ before the merger. Qualitatively the behaviour over the last 4–5 orbits prior to merger was similar for both L12 and L12I as the system rapidly evolved to merger. The main differences can be seen in the binary separation, orbital angular momentum, and binary mass ratio. All of these early differences can be attributed to the accumulation of numerical errors and diffusion. Later in the evolution, when the unstable mass transfer is large and growing exponentially during the few orbits before the merger, the differences between L12 and L12I become less significant since advection dominates in both frames. The main problem with evolutions computed on the inertial frame is that numerical viscosity artificially speeds up the merger. If one is interested in marginally stable/unstable cases, the evolution should be computed in the rotating (comoving) frame.

4.2.4 Non-driven simulations

Most previous merger simulations of WD–WD systems have used some kind of driving mechanism to expedite the merger. Moreover, studies like Motl et al. (2017) and Diehl et al. (2021a) have shown that a shorter driving phase results in a longer evolution to a merger. For practical reasons, the numerical driving rates used in these simulations exceed realistic angular momentum loss rates by factors $\sim 10^9$ or more (from mechanisms like magnetic braking or gravitational waves emission). In this paper, we wanted to investigate the feasibility of simulations closer to realistic driving rates by simply simulating a non-driven system at different resolutions. We know that the lowest mass transfer rate resolvable by a simulation depends on the level of refinement or spatial resolution. By simulating the evolution of the same binary at increasing resolution we may be able to see convergence to a more realistic evolution. Unfortunately, as discussed below, it turned out that the cumulative effect of numerical errors made these simulation results unreliable.

We begin by discussing the low-resolution non-driven run, L10ND. This simulation merged in 37 orbits, even later than L11, which was driven for 1.3 orbits, merging in 25 orbits, and later than L13, which was driven for 2.3 orbits, merging in 10 orbits. However, L10ND merged before L12, which was driven for 1.3 orbits, merging in 39 orbits; see Table 1. This is because, with the limited resolution of L10ND, the initial mass transfer, even without driving, quickly settles on values of the order of $\times 10^{-3} (M_T/P_0)$, higher by a factor of $\sim 2 - 5$ than the mass transfer rate at the start of simulation L12 (middle left panel of Fig. 11). Thus, as we anticipated, the resolution plays a critical role in the early evolution of the mass transfer, and L10ND is probably a poor approximation to a realistic evolution.

We have therefore tried to run non-driven simulations of higher resolution, with 11 and 12 levels of refinement. However, these simulations, which start with lower mass transfer rates than their equivalent driven simulations and thus evolve slower, suffer from some numerical issues, happening after the merging time of their equivalent driven simulations. Specifically, we have found out that in these simulations the merger occurs only because the accretor starts expanding for reasons other than the mass transfer. This expansion of the accretor stems from internal, convection-like, instabilities that form inside the accretor and which were not observed for any of

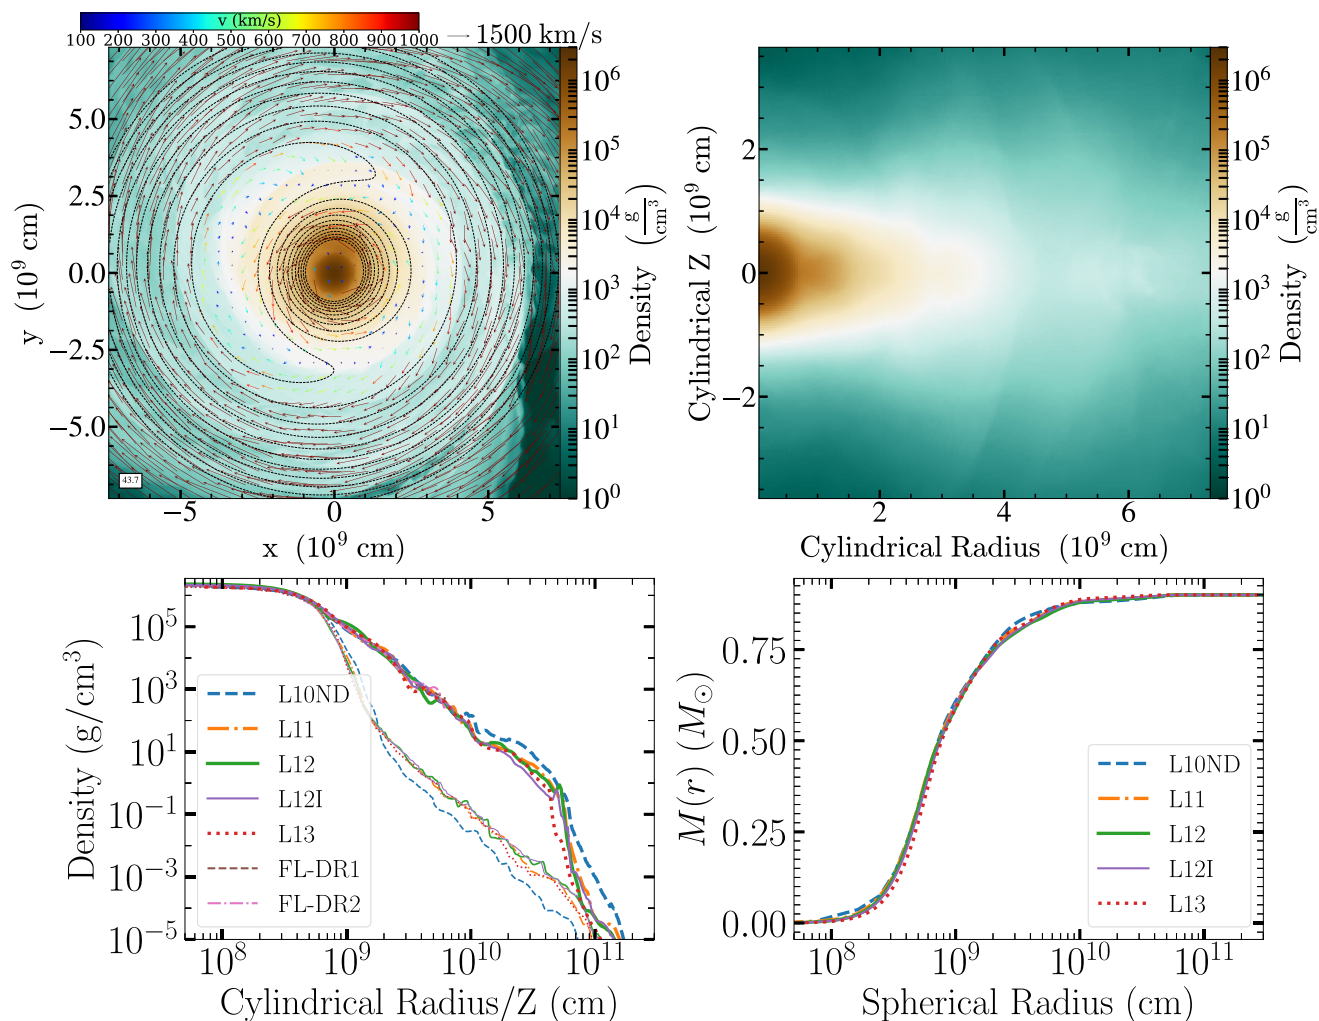


Figure 12. Mass distribution of the merged object five initial orbits past merger. Top row: a density slice along the equatorial plane (left panel), and a mass-averaged over the azimuthal angle density profile (right panel) for simulation L12. Velocity arrows are overlaid with their colour indicating the magnitude of the projected velocities in the rotating frame (see the text and colourbar values as in Fig. 3). Bottom row: the azimuthally averaged density profile (left panel, thick lines), the density profile along the z-axis (left panel, thin lines), and the cumulative mass $M_{\text{in}}(r)$ profile of the merged object for all simulations (right panel). Colours identify the different resolutions according to the legend.

the other simulations (listed in Table 1). We could even reproduce those instabilities in single-star simulations and have linked them to the implemented ZTWD EoS (see Appendix D for more details). In contrast, the simulations of a polytropic binary WD with the same mass ratio of 0.7, evolved with an ideal gas EoS (Diehl et al. 2021b; see Table 1) at four resolutions, all driven with $T_{\text{drv}} = 2.0$, clearly show convergence with respect to merging time. This further supports that the anomalous behaviour observed in some of our non-driven simulations is linked to the EoS. A further investigation of this phenomenon, including the testing of the Helmholtz EoS (Timmes & Swesty 2000) instead of the ZTWD should be considered in order to resolve this issue. We conclude that at this point, with the current implementation of the EoS in OCTO-TIGER, a convergence study of the evolution of a non-driven system is not feasible.

5 THE MERGED OBJECT

We base this discussion on our reference simulation L12, which is performed in the rotating frame, and has 12 levels of refinement. When necessary, we compare with the results of other simulations.

In Section 5.1 we discuss the structure of the merged object, splitting it for convenience into several substructures with characteristic densities and rotational properties, and consider their origins. In Section 5.2 we consider nuclear burning and the potential impact of the corresponding energy deposition. Finally, in Section 5.3 we estimate the dredge-up of core material by extrapolation of results at different resolutions, and discuss the consequences of a small amount of hydrogen in the donor, and its impact on the observed $^{16}\text{O}/^{18}\text{O}$ ratio.

5.1 Structural and rotational properties

In Fig. 12 we focus on the mass distribution of the merged object, five orbits after the merger. In the top row, we show a density slice at the orbital, xy , plane (left), and a density profile mass-averaged over the azimuthal angle (right), both for simulation L12. Cylindrical radius and cylindrical z are measured from the centre of the merged object, i.e. the point of maximum density. Although difficult to see in the top left panel because of the velocity arrows, there is a higher density ‘blob’ at the noon position at radii $\sim 1.2 \times 10^9$ cm, coincident with the cooler blob, clearly visible in Fig. 13, which will be discussed




Figure 13. The SoF five initial orbits past merger. Equatorial temperature slice (left) and a mass-averaged over the azimuthal angle temperature profile (right) for simulation L12.

in the next section. This is a remnant of the core of the donor, which cannot be tidally disrupted by the accretor before it becomes supported by pressure and rotation, but is expected eventually to be sheared by differential rotation. The effect of this blob is also visible in the top right panel, despite the azimuthal averaging. In the bottom left frame, we show the azimuthally mass-averaged density profiles at $z = 0$ (thick lines), and the vertical density profile along the z -axis (thin lines) for all of our OCTO-TIGER simulations. In the bottom right frame, we present the cumulative mass profiles inside a sphere of radius r for all of our OCTO-TIGER simulations. In addition, we plot the radial density structures on the equatorial plane of the FLOW-ER simulations (bottom left; thick lines).

First, we see a remarkably similar structure across all of our simulations, and the following description is valid for each of them. The merged object's radius, defined as the cylindrical radius where the cumulative mass plateaus (bottom right frame), is approximately $R_{\text{out}} \approx 3 \times 10^{10}$ cm. There is some low-density material beyond this radius gradually transitioning to the floor density. A side view (top right frame) and the profiles along the cylindrical radius and cylindrical z -coordinate (bottom-left frame) show that the densest structures consist of a compact oblate spheroid, with an approximate cylindrical radius of $R_{\text{sph}} = 10^9$ cm, and a flattened disc with an approximate outer radius of $R_{\text{disk}} = 3 \times 10^9$ cm around it. We additionally observe that the density contours flare slightly at densities below 10^3 g cm $^{-3}$. Gas velocities in the perpendicular plane are typically smaller than 400 km s $^{-1}$ while the azimuthal velocities on the orbital plane are much larger and as we show later are nearly Keplerian.

Additional description of the mass distribution in the inner regions, valid for *all* the simulations as well, can be gleaned from the bottom left frame of Fig. 12. The nearly constant density $\rho \sim 10^6$ g cm $^{-3}$ spherical region $r \leq r_1 = 5 \times 10^8$ cm corresponds roughly to the core of the CO WD. There is a gradual transition between this core and R_{sph} . For $R_{\text{sph}} \leq R \leq R_{\text{out}}$ cm, the equatorial density in the disc-like or extremely oblate extension beyond the spheroid, falls off as $\rho \propto R^{-3}$ and beyond R_{out} it cuts off precipitously as $\rho \propto R^{-8}$. This is consistent with the $M(r)$ behaviour shown on the bottom right panel.

Furthermore, we can learn the merged object rotation profile (on the equatorial plane) from the velocity arrows in the top left frame. The plotted velocities are in the frame rotating with the grid at a

constant angular frequency of $\Omega_0 = 2\pi/P_0$. Therefore, a velocity v' in the rotating frame transforms to an inertial frame velocity of $v_{\text{in}} = \Omega_0 r + v' = 550$ km s $^{-1} \cdot r_9 + v'$, where r_9 is the cylindrical radius in units of 10^9 cm. Consequently, the velocity at the transition from the spheroid to the disc just outside R_{sph} , on the order of 1000 km s $^{-1}$ in a counterclockwise direction in the rotating frame (seen as red arrows), is ~ 1500 km s $^{-1}$ in the inertial frame, comparable to the sound speed at the SoF. Beyond R_{sph} , v' decreases to near zero at co-rotation (seen as a white ring) located approximately at R_{disk} , and the direction of rotation in the rotating frame inverts beyond corotation, such that farther out, at a radial distance of 4×10^9 cm the rotation is of the order of $v' \approx -1000$ km s $^{-1}$ in the clockwise direction, still turning counterclockwise at ≈ 1000 km s $^{-1}$ in the inertial frame.

In Fig. 14 we show the azimuthally mass-averaged angular velocity profile of the merged object at $z = 0$ for simulation L12, as a function of cylindrical radius five orbits past merger. The angular velocity here is $\Omega = (xv_y - yv_x)/R^2$, where x , y , R , v_x , and v_y are measured with respect to the position and velocity of the merged' object centre.

The angular velocity increases with decreasing radius until just inside R_{sph} , where it starts to decrease again and smoothly matches the solid body rotation of the CO WD near r_1 . The accretor core was initially synchronized to the initial binary frequency, and is still rotating as a solid body at a slightly higher frequency corresponding to the orbital frequency after driving (see Fig. 10). Material inside this radius must be pressure supported since it is well below the local Keplerian velocity, defined as the angular velocity required for pure rotational support, $\Omega_K(r) = \sqrt{(1/R)d\Phi_g/dR}$, where Φ_g is the gravitational potential. On large distances the inner mass becomes nearly a constant and $\Omega_K(r) = \sqrt{GM/R^3}$. On the same figure we plot $R^{-3/2}$ relation normalized to the total mass of $M_T = 0.9 M_\odot$.

The left panel of Fig. 14 additionally shows that the strongest shear occurs between r_1 and R_{sph} . The angular velocity rises rapidly to a maximum value just inside R_{sph} to a sub-Keplerian value consistent with the vertical extent of the spheroid. Beyond the maximum, Ω decreases, but gradually the relative contribution of pressure to supporting the structure declines. The right panel shows the ratio between the local Ω and the local Keplerian value. In all runs, the disc is about 0.8 of the Keplerian value between $r \approx 2R_{\text{sph}} = 2 \times 10^9$




Figure 14. Angular velocity distribution five initial orbits past merger. Left: Gas angular velocity ('X's) and Keplerian angular velocity as computed from the gravitational potential (circles) as a function of distance from the merged object centre for L12. An $\Omega = \sqrt{GM_T/r^3}$ relation expected at large distance is overplotted as a thin line. Right panel: The gas angular velocity divided by the Keplerian angular velocity as a function of radius from the centre of the merged object of all of our OCTO-TIGER simulations.




Figure 15. Specific inertial frame mechanical energy (kinetic plus potential) for all cells for simulation L12 versus specific z angular momentum with respect to the centre of mass at four times: $t = 0, T_{\text{merge}} - 5, T_{\text{merge}}$, and $T_{\text{merge}} + 5$. The colour bar on the right of each frame shows the mass in each bin. See the text for a detailed description.




Figure 16. The azimuthally averaged radial profile of temperature on the orbital plane (thick lines) and the vertical temperature distribution along the positive z -axis for simulation L12, five initial orbits past merger.

cm and $R \approx 10^{10}$ cm. Beyond this radius the disc is probably not in equilibrium yet. The material that is still falling back after tidal disruption is unlikely to be supported by pressure, and therefore must have sub-Keplerian azimuthal velocities.

Fig. 15 depicts the distribution of specific inertial-frame mechanical energy as a function of the specific z angular momentum with respect to the centre of mass of the system in the L12 simulation at four selected times. We divide the specific energy and the specific angular momentum in the range shown in the figure to 128×128 equally spaced bins. The colour indicates the mass in each bin. This provides a picture of the dominant mass structures at different stages of evolution. A similar technique was used by Hayashi et al. (2021) to describe the tidal disruption of a neutron star by a black hole.

At $t = 0$, all the mass is in the binary components, with the minima corresponding to the central densest and most bound element in each component. The specific angular momenta of the centres of the components are in the ratio q^2 , with the donor having the largest central $l_z = (M_1/M_T)^2 a_0^2 \Omega_0 = 2.2 \times 10^{17}$ cm² s⁻¹. The kinetic energy is due just to the synchronous rotation of the system at the initial orbital frequency Ω_0 . With R the cylindrical radial distance of any element from the binary axis, we have

$$E = \frac{1}{2}(\Omega_0 R)^2 + \Phi_g = \Omega_0 l_z + \Phi_{\text{eff}},$$

where $\Phi_{\text{eff}} = \Phi_g - \frac{1}{2}(\Omega_0 R)^2$ is the effective potential. The surfaces of both accretor and donor are $\Phi_{\text{eff}} = \text{Const}$ surfaces, so E is linear on l_z with slope Ω_0 . Two lines with that slope are shown for reference.

At $t = T_{\text{merge}} - 5P_0$, both components are easily distinguishable, but both centres have moved to lower l_z and more negative (bound) values of E . The stream material fills the region at low l_z between the two components. The gas leaking out of L_2 appears marginally bound/unbound at large l_z . At $t = T_{\text{merge}}$, the centre of the accretor material has moved very close to the centre of mass of the system at $l_z = 0$, and the donor is being disrupted, so its binding energy is significantly lowered.

Finally, at $t = T_{\text{merge}} + 5P_0$, the accretor and the centre of mass have nearly converged to $l_z = 0$, and the binding energy is even greater. The material at $l_z \leq 0.5 \times 10^{18}$ cm² s⁻¹ shown in yellow tones is the most tightly bound accretor matter, rotating nearly as a solid body $l_z = \Omega R^2$, where Ω is slightly faster than the initial value (see left panel of Fig. 14). For reference, we have drawn lines

following $|E| \propto l_z^{-2}$ corresponding to Keplerian orbits around point masses 0.5, 0.6, 0.7, 0.8, and, $0.9M_\odot$. For $l_z \geq 0.4 \times 10^{18}$ cm² s⁻¹ the highest density structure is consistent with Keplerian or sub-Keplerian orbits around a mass $0.9M_\odot$. At intermediate values of l_z , we find structures, seen in green and yellow tones, corresponding to a boundary layer and a vestige of the donor core (the blob) still visible at $E = -0.75 \times 10^{17}$ erg g⁻¹, $l_z = 0.2 \times 10^{18}$ cm² s⁻¹. See again Fig. 14 for comparison. Note also the green-yellow feature at approximately 0.9×10^{18} cm² s⁻¹ $\leq l_z \leq 1.1 \times 10^{18}$ cm² s⁻¹, corresponding to material leaked out of L_2 , whose highest l_z portion is slightly unbound. All of the above results are consistent with the donor WD being partially disrupted by tides, leaving the core of the donor still surviving, and the envelope being stripped and sheared, with a small amount being leaked out of L_2 .

5.2 Nuclear fusion in the merged object

During a merger of two WDs, temperatures and densities are such that nuclear fusion of helium into carbon and oxygen is taking place at a shell of accreted donor material around the accretor, termed the shell of fire (SoF). In Fig. 13 we show a temperature slice at the orbital, xy , plane (left), and a temperature profile mass-averaged over the azimuthal angle (right), both for simulation L12. As in Fig. 12, cylindrical radius and cylindrical z are measured from the centre of the merged object, the point of maximum density. The SoF is evident as the hot blue ring around a central yellow circle. In Fig. 16 we present an azimuthally mass-averaged temperature profile at $z = 0$ (thick line), and a temperature profile along the positive z -axis (thin line) for simulation L12. All of our OCTO-TIGER simulations show a similar temperature profile with a peak temperature between $1 - 2 \times 10^8$ and slightly hotter temperatures at the poles. Additionally, due to their lower resolution, in simulations L10ND and L11, the spherical central core possesses higher temperatures of 2×10^7 K and 4×10^6 K, respectively.

The presence of a cooler remnant of the donor’s core, which has not had time to shear and to mix with the surrounding gas, lowers the azimuthal average temperature on the orbital plane (Fig. 16, thick line). As a result, the highest temperatures of the SoF occur at higher latitudes. This cooler remnant can be seen clearly in the equatorial slice of Fig. 13 as a red ‘blob’ adjacent to the SoF at positive y direction, and in the averaged profile (right panel of Fig. 13) as a yellow feature at $R_{\text{sph}} < R < R_{\text{disk}}$, which extends symmetrically from $z = -0.4 \times 10^9$ cm to $z = 0.4 \times 10^9$ cm along the vertical axis. The presence of this cooler blob had already been noticed in previous simulations (Staff et al. 2012, 2018).

A more useful display of the He-burning zones is gleaned from Fig. 17, where we show the time evolution of the mass distribution in density and temperature around the time of the merger. To construct this mapping we divide the domains of density and temperature into evenly spaced 128×128 logarithmic bins, and we depict using a colour bar how much mass is in each bin. Consequently this display reveals also the main structures of the flow at the three selected times. The structures that extend above the black line meet the conditions for triple alpha burning. Focusing on the bins with most mass, at $t = 33.72P_0$, near the bottom right corner of the left panel, we can see two structures superimposed and roughly shaped as paint-brushed European-style ‘1’s, the accretor core as a nearly vertical blue band extending to $\rho \sim 10^6$ g cc⁻¹, with the accretor envelope represented by the serif extending down and to the left. Similarly, the donor core appears as a vertical band at $\rho \sim 2 \times 10^5$ g cc⁻¹, with its envelope as the corresponding serif. The ‘knee’-




Figure 17. Density–temperature phase plots for L12 at three times corresponding (from left to right) to $T_{\text{merge}} - 5P_0$, T_{merge} , and $T_{\text{merge}} + 5P_0$, respectively. The black line in each panel denotes the threshold for triple alpha burning. The temperature and density are divided each to 128 bins, and the colour indicates the mass in each bin.

like feature above the 3α boundary and its downward extension joining the accretor is the accretion belt with the SoF. At $t = 38.72P_0$ the knee has moved to higher temperatures and densities, and the total mass above the 3α boundary has increased. The donor is still discernible. Finally, five orbits after the merger, the donor is gone, the hot accretion belt and disc are visible as a broad fanning feature with an approximate slope of $2/3$, which culminates on a slightly cooler knee (than at time $t = 38.72P_0$) above the 3α boundary at even higher densities and extends downward to meet the accretor core. There is now more mass above the 3α boundary consistent with Fig. 18 (see below).

In Fig. 18 we show, on the top row, the distribution of He-burning power on the equatorial xy plane (left) and on the vertical xz plane (right). Overplotted are four contours of densities (ordered spatially from the exterior to the interior of the merged object) of 10^3 , 10^4 , 10^5 , 10^6 g cm $^{-3}$, in yellow, and two contours of temperatures, one of 5×10^7 K in red, and a second of 10^8 K in purple. On the bottom left frame, we display the mass of the SoF as a function of time, where the mass of the SoF is determined by summing the mass in the bins above the 3α threshold in Fig. 17. This is an improved method to measure the SoF mass as compared to what was done in earlier simulations (Staff et al. 2012, 2018) since we find that some helium burning can still take place at densities of $\sim 10^4$ g cm $^{-3}$ (mainly near the poles) or alternatively at temperatures less than 10^8 K (e.g. at the equator). In the bottom right panel of the same figure we plot the total luminosity generated by the triple alpha burning as a function of time in our OCTO-TIGER simulations.

Comparing the results for different resolutions, it is clear that a higher resolution is needed to get an accurate estimate of the mass of the burning shells and the energy released. In the lower resolution runs, in which the accretor is resolved across fewer cells, the shells interior to the inner boundary of the SoF become hotter with time, so the SoF effectively expands inwards to higher densities and the amount of mass in the SoF quickly grows and therefore is probably

being overestimated. For higher resolutions the maximal density that burns increases more slowly with time and the SoF mass is smaller as a result.

We list in Table 2 representative values of the SoF for simulations L10ND, L11, and L12 with the purpose of examining the convergence of these properties with respect to resolution. We do not include in this analysis simulation L12I because unlike the rest of the other simulations it had been performed in the inertial frame of reference and not in the rotating frame. We also do not include simulation L13 because it shows slight peculiar features in its internal structure like a small decrease in the accretor’s central density (bottom left panel of Fig. 11), and a high central rotation (right panel of Fig. 14). In the second column of Table 2, we list a representative value for the temperature in the SoF, T_{SoF} , that we define as the temperature in which the nuclear power in each simulation is the strongest. This temperature does not vary much with resolution. It oscillates shortly after the merger between $1.3 - 2.3 \times 10^8$ K, but then settles down. Averaging over a time period that starts at one initial orbital period after the merger we obtain $T_{\text{SoF}} = 1.6 \times 10^8$ K.

The SoF density, ρ_{SoF} (third column) is calculated as the maximal density in the SoF averaged over the same time period, which decreases for higher resolution. We then perform a Richardson extrapolation (Richardson & Gaunt 1927), by taking results for three increasing resolutions, and extrapolate to the continuum limit to obtain an estimate for the case for cell size going to zero. This yields a density value of 2.9×10^5 g cm $^{-3}$.

To check convergence for the mass of the burning shells, we average the values of M_{SoF} (as shown in the bottom left panel of Fig. 18) at two time points: $t = T_{\text{merge}} + P_0$, and the latest simulated time. The obtained values are plotted as three horizontal lines (dashed blue, dashed-dotted orange, and thick solid green, for simulations L10ND, L11, and L12, respectively) in the same figure and are also listed in the fourth column of Table 2. The SoF mass shows a second-order convergence to a value of $M_{\text{SoF}} \simeq 0.08$ with respect to




Figure 18. Upper row: Equatorial, xy (left) and meridional, xz (right) slices of helium burning power density five initial orbits past merger for L12. Bottom row: The gas mass that has the conditions to burn helium (left) and its burning power (right) as a function of time. The three horizontal lines, dashed blue, dashed-dotted orange, and thick solid green, on the bottom left panel, denote a burning mass averaged between one orbital period after the merger and the latest simulated time, for simulations L10ND, L11, and L12, respectively. These averages are used to derive the continuum limit of our resolution study (plotted as the lowest horizontal dashed black line; see Table 2 as well).

Table 2. Convergence with increasing resolution of SoF properties.

Simulation	T_{SOF} (MK)	ρ_{SOF} (g cm^{-3})	M_{SOF} (M_{\odot})	$M_{\text{SOF}}^{16\text{O}}$ (M_{\odot})	f_{core}
L10ND	160	7.0×10^5	0.244	0.062	0.51
L11	160	3.3×10^5	0.128	0.032	0.50
L12	160	3.0×10^5	0.094	0.019	0.40
Continuum limit	160	2.9×10^5	0.081	0.01	0.25
Convergence order	–	3.4	1.8	1.2	–

Note. We perform a Richardson extrapolation (Richardson & Gaunt 1927) to obtain the continuum limit for ρ_{SOF} , M_{SOF} , and $M_{\text{SOF}}^{16\text{O}}$. $f_{\text{core}} = 2M_{\text{SOF}}^{16\text{O}}/M_{\text{SOF}}$.

resolution (see the lowest horizontal dashed black line in the same figure and also Table 2).

The amount of mass in the SoF and its composition (as we will discuss in the next section) greatly affect the final byproducts of the

partial helium burning. These burning products, later on, as the post-merger object expands, cools, and evolves to become a carbon star, will get mix with the envelope due to convection. Therefore, they will also determine the surface abundance at the RCB phase. Our aim is to infer from the simulations the estimated conditions expected due to a dynamical merger of a $q = 0.7$ system, from which we can ultimately estimate the surface values as the burning products mixed with the envelope.

Ideally one would like to include the energy deposition by nuclear reactions as the 3D simulation proceeds and to calculate the dynamical reaction to the deposition of this extra energy. However, OCTO-TIGER does not yet include a nuclear reaction network and therefore we do not include in the simulation the effects of nuclear energy generation. These reactions probably play a role in reshaping the merged object on a dynamical time-scale, and in changing the abundances on a nuclear time-scale. What we can say is that the total energy generated during the simulation can be estimated as the

peak luminosity times $\sim 10P_0$, yielding $E_{3\alpha} \approx 4 \times 10^{44}$ erg, which is negligible compared to the various energy components of the system (see Fig. 8).

Moreover, we find that the highest power density of the nuclear energy is generated at gas with temperatures $T = 1.4 - 2.0 \times 10^8$ K, and densities $\rho = 1 - 2.5 \times 10^5$ g cm $^{-3}$, and equals roughly to $\epsilon_{3\alpha}^{\max} = 5 \times 10^{12}$ erg cm $^{-3}$ s $^{-1}$. At this range of temperatures and densities, the minimal thermal energy density, assuming a composition of ionized carbon–oxygen, equals $E_{\text{th, min}} \approx 8 \times 10^{20}$. Therefore $\approx 10^7$ s are required (roughly 10^5 orbits) for the nuclear reaction to become important at these regions, which is beyond the scope of the hydrodynamic simulation, and thus can be neglected.

The merged object at the end of our simulations, 5 or even 10 initial orbits after the merger, is not yet spherically (or even fully axially) symmetric. However, following the long-term evolution of the post-merger object by means of 3D hydrodynamic simulations is not feasible as the time-step is ultimately limited by the Courant condition. Mapping the 3D object into a 1D implicit code such as MESA (Paxton et al. 2011, 2013, 2015, 2018, 2019) is a possible way of tackling the long-term evolution provided that we can find an adequate method of averaging the 3D structure of the merged object, which preserves the object’s mass and its angular momentum distribution.

In this paper, we average the merged object along different angles based on a similar procedure to that described in section 3.1 of Munson et al. (2021). Cells between a polar angle $\theta + d\theta$ and $\theta - d\theta$ are averaged around the azimuthal angle to obtain the quantity at a given stellar radius and the averaging is done for conical shells along both the positive and negative z-axis. The results are consistent with the temperature distribution shown in Fig. 16 and the burning power distribution shown in Fig. 18. By comparing these results to the models of Munson et al. (2021), we can evaluate what can occur after the merger event. Since the models reach He-burning temperatures during the merger, assuming they will follow an RCB-like evolution is reasonable. The input luminosity from the steady He-burning shell should expand the envelope and bring the surface luminosity from 4000 to 10 000 L_{\odot} (Menon et al. 2013, 2019; Lauer et al. 2019; Crawford et al. 2020).

5.3 The dredge-up of ^{16}O and the inferred $^{16}\text{O}/^{18}\text{O}$

The dredge-up of ^{16}O , its impact on the surface ratio of $^{16}\text{O}/^{18}\text{O}$ (the oxygen ratio), and the discrepancies between the estimated dredge-up by SPH and grid codes were already discussed in Staff et al. (2018). Here, we re-examine the amount of dredged-up core material using a similar approach as in Staff et al. (2018), at several levels of resolution, using an updated OCTO-TIGER, and an improved definition of the SoF. Using the results obtained at three different and increasing resolutions, we estimate the continuum limit dredge-up amount using the Richardson extrapolation.

We are not aiming here to obtain an accurate estimate of the oxygen ratio since that would require following a nuclear reaction network simultaneously with the hydrodynamics of the merger and the precise timing of when the outer convective zone connects with the burning products of the SoF (Menon et al. 2013; Crawford et al. 2020; Munson et al. 2021). Instead, we assume constant conditions, that is constant temperature and density, in the regions of helium burning, and inspect only the most relevant reactions to the production of ^{16}O and ^{18}O to infer a minimal possible value of $^{16}\text{O}/^{18}\text{O}$ in these regions. We also assume that the reactions proceed to completion which should occur at a time-scale dictated by the reaction with the

slowest rate. The temperature, density, and initial composition of this analysis are derived from our hydrodynamic simulations. A similar single-zone nucleosynthesis approach have been performed in Staff et al. (2012), where in that paper the analysis has been carried-out with a nuclear reaction code. We build on their results in diagnosing the most relevant processes to yield a simple analytic estimate without running a nuclear reaction code. Our main conclusion is that the dredge-up of ^{16}O is accompanied by ^{12}C , which can help the production of ^{18}O , if enough hydrogen is present in the donor envelope.

As we showed in Section 5.2, an accurate estimate of the *total* mass of the SoF, and regardless of composition, at any time for a given resolution, is obtained by adding the mass above the threshold for triple alpha burning in the $T - \rho$ distribution (e.g. the black line in Fig. 17 for selected times during the L12 run). The results of these estimates for all runs as a function of time are shown in the bottom left panel of Fig. 18.

In Fig. 19 we show the mass of dredged up ^{16}O (and also ^{12}C , assuming equal masses in the dredged-up accretor material) as a function of time for all of our simulations. Since we evolved three distinct fluids in our simulations: one for the accretor’s core (CO), a second for the accretor’s envelope (He), and a third for donor material (also He), we could calculate how much mass of the accretor’s core material, i.e. the gas that was originally located at the accretor’s core, has densities smaller than 10^5 g cm $^{-3}$ and take half for the amount of ^{16}O (Fig. 19 left panel). This is to be compared with fig. 4 from Staff et al. (2018).

We also calculated a more precise measurement of how much mass of the accretor’s core fluid resides in regions where conditions for helium burning (by the triple alpha reaction) exist for several resolutions, and derive the amount of ^{16}O (and ^{12}O), again by taking half. We plot this quantity in the right panel of Fig. 19. The main difference between these quantities is that on the left panel, mass that resides outside the SoF and at lower densities, is still counted, while the right panel informs us the composition in the regions where triple-alpha burning happens. Similar to the bottom left panel of Fig. 18, we additionally plot on the right panel an average $M_{\text{SoF}}^{16\text{O}}$ between two time points, $t = T_{\text{merge}} + P_0$, and the latest simulated time, for simulations L10ND, L11, and L12 as horizontal lines (in dashed blue, dashed-dotted orange, and thick solid green, respectively). These averages are used to derive the continuum limit dredge-up value of our resolution study (plotted as the lowest horizontal dashed black line; see Table 2 as well).

The general behaviour of the dredge-up curves on the left panel of Fig. 19 consists of three phases: an initial fast rise, perhaps dominated by numerical diffusion, a monotonic slow increase until the merger, and a steep increase during the merger, most likely a combination of real dredge-up and numerical diffusion. It makes sense therefore to take values of the dredge-up at the same time after the merger at different resolutions and to apply the Richardson extrapolation to estimate the corresponding value in the continuum limit. The curve thus obtained is shown as a black dashed line in the left panel of Fig. 19 and is composed of ~ 100 data points with an average convergence order of ~ 1 (and a standard deviation of ~ 0.2). Based on this analysis we conclude that the final mass of dredged-up ^{16}O , after subtracting the fast rise ‘shelf’ value shown as horizontal black line at pre-merger times, is $\sim 0.02M_{\odot}$. However, not all of the dredged-up ^{16}O ends up in the SoF, see the discussion below. The reason that analogous SPH simulations essentially found no dredge-up may be attributed to the known inability of standard implementations of SPH to correctly resolve mixing at fluid boundaries unless

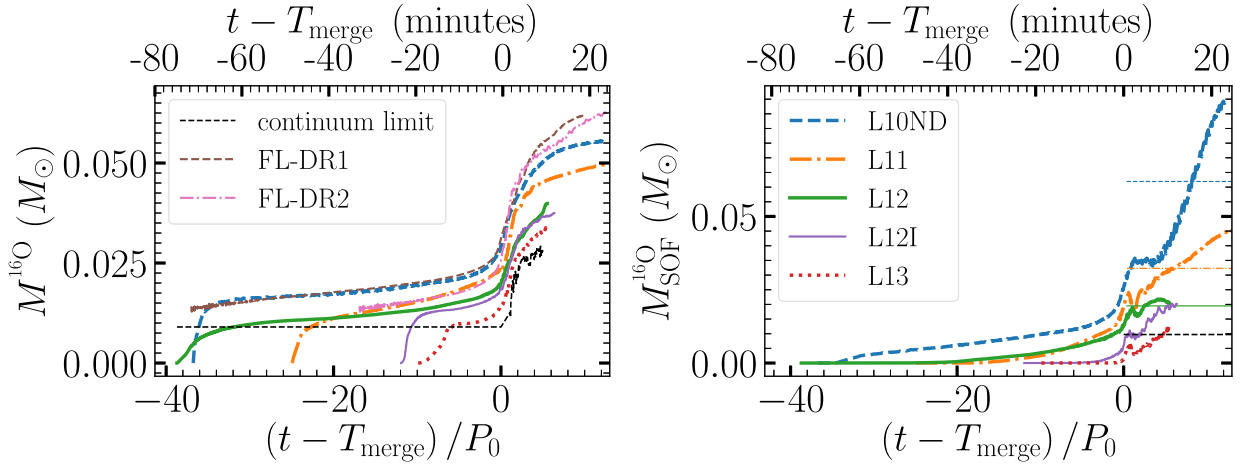


Figure 19. Left: Dredge-up of oxygen-16, defined as mass of oxygen-16 residing at densities lower than 10^5 g cm^{-3} . The dashed black curve labelled ‘continuum limit’ is the Richardson extrapolation for cell size going to zero, derived from simulations L11, L12, and L13. It is composed of ‘plateau or shelf’ value before the merger, and actual values 1–5 orbits after the merger. The ‘shelf’ value is calculated at time $t_{\text{shelf}} = -0.5(T_{\text{merge}} - 4P_0)$. Right: mass of oxygen-16 in the SoF, where the conditions for helium burning exist. The three horizontal lines, dashed blue, dashed-dotted orange, and thick solid green, denote a value averaged between one orbital period after the merger and the latest simulated time, for simulations L10ND, L11, and L12, respectively. These averages are used to derive the continuum limit value of our resolution study (plotted as the lowest dashed black line; see Table 2 as well). The upper axis in both panels shows time in minutes, where the conversion is based on the initial orbital period of simulation L12, $P_0 = 113.6 \text{ s}$.

certain modifications are adopted (Read, Hayfield & Agertz 2010; Ruiz-Bonilla et al. 2022).

The observations suggest that in most RCB stars the ratio $^{16}\text{O}/^{18}\text{O} \leq 10$. Some RCB stars have higher values for this ratio but still much less than the solar value ~ 500 . Most of the ^{18}O comes from the reaction $^{14}\text{N}(\alpha, \gamma)^{18}\text{F}(\beta^+)^{18}\text{O}$. Thus the amount of ^{14}N present in the SoF limits the amount of ^{18}O produced, since some ^{18}O may burn to ^{22}Ne . The amount of ^{14}N in the SoF comes from two possible sources: what was originally present in the He WD, plus what is synthesized in the SoF during burning through proton captures $^{12}\text{C}(p, \gamma)^{13}\text{N}(\beta^+)^{13}\text{C}(p, \gamma)^{14}\text{N}$, which requires the presence of sufficient H (see below).

Considering first the simpler case in which no H is available in the SoF, the mass of ^{18}O would be limited by the small amount of ^{14}N originally present in the donor. The amount of ^{16}O would mainly depend on the dredge-up from the accretor’s core. Assuming an initial ^{14}N mass fraction of $X_{^{14}\text{N}}^0$ in the SoF, and $0.5f_{\text{core}}$ mass fractions of ^{12}C and ^{16}O while the rest is helium, the minimal mass-fraction ratio of $^{16}\text{O}/^{18}\text{O}$ in the SoF would be

$$\left(\frac{^{16}\text{O}}{^{18}\text{O}}\right)_{\text{SoF}} = (7/18)f_{\text{core}}/X_{^{14}\text{N}}^0, \text{ with no hydrogen.} \quad (3)$$

Here, $f_{\text{core}} = 2M_{\text{SoF}}^{16}\text{O}/M_{\text{SoF}}$ is the mass fraction of the accretor’s core fluid in the SoF, which can be inferred from our simulations. In our low resolution simulations, L10ND and L11, we obtain $f_{\text{core}} \simeq 0.5$, while simulation L12 yields $f_{\text{core}} \simeq 0.4$, and the extrapolation to infinite resolution yields a lower fraction of $f_{\text{core}} = 0.25$ (a first-order convergence; see horizontal lines in the right panel of Fig. 19 and Table 2). With no hydrogen and this amount of dredge-up, values of $^{16}\text{O}/^{18}\text{O}$ less than 10 are impossible unless $X_{^{14}\text{N}}^0 \geq 4$ per cent f_{core} . However, the amount of ^{14}N is diluted by the accretor material that has been dredged-up to the SoF, so the mass fraction of ^{14}N in the He WD donor star, $X_{^{14}\text{N}}^{\text{He-WD}}$, should be $X_{^{14}\text{N}}^{\text{He-WD}} \geq 4$ per cent $(f_{\text{core}}/(1 - f_{\text{core}}))$ for $^{16}\text{O}/^{18}\text{O} \leq 10$. This imposes a substantial restriction, specifically because observations of most RCB stars show mostly subsolar metallicities (e.g. Asplund et al. 2000). Obtaining a ratio of the

order of unity without the presence of hydrogen in these cases is unlikely.

With H in the SoF, on the other hand, we can create more ^{14}N by two proton captures for every ^{12}C , as seen in the reaction chain $^{12}\text{C}(p, \gamma)^{13}\text{N}(\beta^+)^{13}\text{C}(p, \gamma)^{14}\text{N}$. Because the half-life of ^{13}N is 9.9 min and the proton capture rate of $^{12}\text{C}(p, \gamma)$ in the SoF is 0.6 s, it is fair to assume that the available ^{12}C will capture protons before $^{13}\text{C}(p, \gamma)^{14}\text{N}$ takes place. Furthermore, because the average temperature in the SoF is only 150MK, there will be a negligible amount of additional ^{12}C generated via triple-alpha processes. Therefore, only the protons remaining after ^{12}C is depleted in the SoF are available for $^{13}\text{C}(p, \gamma)^{14}\text{N}$, and only they will contribute to the formation of additional ^{18}O , which will reduce the $^{16}\text{O}/^{18}\text{O}$ ratio. When there are enough protons to convert all the ^{12}C to ^{14}N , any additional protons do not contribute to the production of additional ^{18}O and will not change this ratio. The $^{16}\text{O}/^{18}\text{O}$ in the SoF can therefore be written as:

$$\left(\frac{^{16}\text{O}}{^{18}\text{O}}\right)_{\text{SoF}} = (7/18)f_{\text{core}}/X_{^{14}\text{N}}^0 \times \begin{cases} 1 & X \leq f_{\text{core}}/24, \\ (1 + f_1 - f_2)^{-1} & f_{\text{core}}/24 \leq X \leq f_{\text{core}}/12, \\ (1 + f_2)^{-1} & f_{\text{core}}/12 \leq X, \end{cases} \quad (4)$$

where $f_1 = 14X/X_{^{14}\text{N}}^0$, $f_2 = 7f_{\text{core}}/(12X_{^{14}\text{N}}^0)$, and X is the hydrogen mass fraction in the SoF.

We plot this function in Fig. 20 for two representative values of $X_{^{14}\text{N}}^0$ of 5×10^{-3} (left panel) and 5×10^{-4} (right panel; corresponding to subsolar metallicities) in solid lines and for four values of f_{core} : 0.01 (blue), 0.25 (orange), 0.4 (green), and 0.5 (red). Later on, as the burning proceeds, some ^{16}O can be created by the remaining ^{13}C via $^{13}\text{C}(\alpha, n)^{16}\text{O}$, which would be contributing to increasing the ratio. The ratio including this reaction is plotted in the same figure as dashed lines.

As can be seen, there are two regimes with a sharp transition between them. For low X values, the ratio is mostly dependent of the initial nitrogen abundance (comparing the left and right panels with $X_{^{14}\text{N}}^0$ of 0.5 per cent, and 0.05 per cent, respectively).




Figure 20. The $^{16}\text{O}/^{18}\text{O}$ ratio in the SoF as a function of H mass fraction, X , assuming constant density and temperature, for two representative initial abundances of ^{14}N . Left: $X_{^{14}\text{N}}^0 = 0.5$ per cent. Right: $X_{^{14}\text{N}}^0 = 0.05$ per cent. Solid lines are according to equation (4). Dashed lines also include alpha captures on ^{13}C , which creates more ^{16}O and therefore increase the ratio. We additionally assume that all the involved reaction chains $^{12}\text{C}(p, \gamma)^{13}\text{N}(\beta^+)^{13}\text{C}(p, \gamma)^{14}\text{N}$, $^{14}\text{N}(\alpha, \gamma)^{18}\text{F}(\beta^+)^{18}\text{O}$ proceed to completion, and that ^{12}C captures all the protons before $^{13}\text{C}(p, \gamma)^{14}\text{N}$ takes place.

For higher X the ratio decreases to a value ≤ 1 . The location of the transition is a function of the accretor mass fraction in the SoF, f_{core} . It can be generally written as $X > f_{\text{core}}/24$, where at twice this mass fraction at $X > f_{\text{core}}/12$, the ratio falls down to < 1 . For our infinite resolution extrapolation, $f_{\text{core}} = 0.25$, a minimal amount of $X > 1$ per cent is required to obtain the observed values in RCBs.

To estimate the H mass fraction in the SoF we assume that any amount of hydrogen resident in the He WD envelope (outside the hydrogen free core) gets mixed completely into the SoF, i.e. $X = M_{\text{He-env}}^{\text{H}}/M_{\text{SOF}}$, where $M_{\text{He-env}}^{\text{H}}$ is the hydrogen mass in the envelope of the He WD. Under this assumption an amount of $M_{\text{He-env}}^{\text{H}} = (f_{\text{core}}/24)M_{\text{SOF}} \approx 10^{-3}M_{\odot}$ is required. For the case considered here, where the He WD is the result of a previous complex binary interaction, which may include several episodes of mass transfer and one or two full common envelope interactions, the amount of hydrogen in the He WD envelope can only be poorly constrained. Alternatively, if a merger takes place in an environment that is richer in H, like a merger inside a common envelope or a merger with a sub-giant star as proposed by Shen, Blouin & Breivik (2023), they could also produce the right ratios. In the latter cases the rest of the hydrogen envelope must become unbound in the process, so by the time the post-merger object evolves to its RCB phase (on the order of $\sim 10^2 - 10^3$ yr), no hydrogen is seen.

Finally, two other factors can increase the $^{16}\text{O}/^{18}\text{O}$ ratio, although we claim here that they play only minor roles. One is the amount of ^{16}O that is dredged-up from the accretor *outside* the SoF. If we assume a total mixing of the mass in the SoF to the surface, the surface ratio will be

$$\begin{aligned} (^{16}\text{O}/^{18}\text{O})_{\text{surf}} &= \left(M_{\text{SOF}}^{16\text{O}} + M_{\text{out}}^{16\text{O}} \right) / M^{18\text{O}} \\ &= (^{16}\text{O}/^{18}\text{O})_{\text{SOF}} \left(1 + M_{\text{out}}^{16\text{O}} / M_{\text{SOF}}^{16\text{O}} \right), \end{aligned} \quad (5)$$

where $M_{\text{out}}^{16\text{O}}$ is the ^{16}O mass outside the SoF as a result of the dredge-up from the accretor. We find in our infinite resolution extrapolation $0.02 M_{\odot}$ of ^{16}O at low densities (Fig. 19 left panel)

and that $0.01 M_{\odot}$ of it resides in the SoF (Fig. 19 right panel), while the remaining $0.01 M_{\odot}$ must reside outside the SoF. Consequently $M_{\text{out}}^{16\text{O}} \simeq M_{\text{SOF}}^{16\text{O}}$, and this dredge-up to outer layers only doubles the ratio. Additionally, the ^{18}O will eventually burn to ^{22}Ne through the reaction $^{18}\text{O}(\alpha, \gamma)^{22}\text{Ne}$. However, Staff et al. (2012) found that the relatively low temperature in the SoF will delay this reaction to $10^{12} \text{ s} \approx 30,000 \text{ yr}$ (see the green line in their fig. 16, for example). At this point, the convection deep into the SoF has already stopped (see e.g. Menon et al. 2013; Crawford et al. 2020; Munson et al. 2021), and therefore the $^{16}\text{O}/^{18}\text{O}$ ratio in the SoF will no longer affect the surface ratio. In Munson et al. (2021) the convection to the SoF persists up until 100 yr and then stops. This gives a window of time for the production of ^{18}O without being processed further to ^{22}Ne .

6 SUMMARY AND CONCLUSIONS

In this paper, we have carried out a set of simulations of WD mergers to further our understanding of the origin and phenomenology of RCB stars. All our simulations have total masses in the narrow range $0.85 - 0.9M_{\odot}$ because of observational constraints, and a mass ratio $q = M_2/M_1 = M_{\text{donor}}/M_{\text{accretor}} = 0.7$ because previous simulations have indicated that this ratio yields the optimal conditions for incomplete helium burning. Our simulations extend previous studies by using a new and improved code OCTO-TIGER and comparing the outcomes for a range of resolutions, at 10, 11, 12, and 13 levels of refinement. The finest grid cell for L10ND was $1.2 \times 10^{-3} R_{\odot}$, while the finest cell for L13 was $1.43 \times 10^{-4} R_{\odot}$ (Table 1). The outcomes of our resolution study show that a finer grid results in better angular momentum conservation (Fig. 2) and allows us to follow the evolution with lower mass transfer rates (Fig. 11). However, this also translates to a longer and costlier evolution to merger. We find excellent agreement between our L13 evolution and the analytic predictions right up to the tidal disruption and merger, where the analytic treatment fails (Appendix B). This comparison serves both to verify our method and illustrates the limitations of the simple analytic approach.

Another comparison was with respect to the choice of frame (Section 4.2.3). Inertial frame simulations suffer from numerical diffusion due to the advection of the binary stars across the grid cells, which results in higher initial mass transfer rates and a rapid evolution through a merger (Fig. 11). Rotating frame simulations avoid these problems and thus they are much more suitable for the simulation of marginally stable or unstable binaries.

The comparison of our simulations with equivalent ones carried out with a different grid code, FLOW-ER, have yielded remarkably similar results (Appendix C). However, OCTO-TIGER has the advantage of using adaptive mesh techniques, allowing us to simulate a larger domain than in FLOW-ER. In OCTO-TIGER simulations we therefore could follow the merged object up to large distances and catch that the density profile outside the disc falls like $\rho \propto R^{-8}$ (Fig. 12).

We then analysed the merged object in Section 5. Its structure is an inner pressure-supported, sphere, slightly oblate, rotating as a solid body, with approximate radius $R_{\text{sph}} = 1 \times 10^9$ cm surrounded by a disc (Fig. 12). The inner part of the disc is in Keplerian rotation, while the outer parts rotate slower than the Keplerian speed. This could be due to partial pressure support or that these outer regions would eventually fall back (Fig. 14).

We have analysed the temperature and density conditions of the gas after the merger to determine how much of the gas is likely to undergo nuclear fusion (Section 5.2). As in previous studies, we have found a region of hot temperatures around the accretor with the right conditions for helium burning. However, due to the presence of a cold blob of gas, a remnant of the disrupted donor, regions of higher latitudes are hotter than the equatorial, and the SoF does not consist of complete shell, but has some hole(s) in it (Figs 13, 16, and 18). The amount of fusing gas is a function of resolution, but it converges to $\sim 0.08 M_{\odot}$ (see Fig. 18 and Table 2). The energy generated by the nuclear reactions over the time span of ~ 20 min is $\sim 10^{44}$ erg still negligible compared to any other energy component in the simulation, thus justifying the fact that we do not include nuclear effects on the evolution of the simulation.

The most important goal of our research was to re-examine the dredge-up of material from the accretor at several resolutions and to obtain an improved estimate of the dredged-up mass by extrapolating to an ‘infinite’ resolution (Fig. 19). Our extrapolated value of $M^{16\text{O}} = 0.02 M_{\odot}$ at lower densities than 10^5 g cm^{-3} (Fig. 19 left panel) bridges the gap in the amount of dredged-up ^{16}O mass found in previous $q = 0.7$ simulations between SPH versus finite-volume codes (see table 1 of Staff et al. 2018), and our result sits right at the middle. In addition, the extrapolated amount of $0.02 M_{\odot}$ of accretor mass in the SoF (Fig. 19 right panel and Table 2), and outside the SoF, is lower than previously found (see Staff et al. 2012 table 3), which consequently results in two times smaller accretor mass fraction in the SoF of $f_{\text{core}} = 0.25$ (see Table 2 versus Staff et al. 2012 table 2). However, our analysis based on the most relevant nuclear reactions (Fig. 20) indicates that even with this relatively mild dredge-up of ^{16}O from the accretor, creating ^{18}O only from the available ^{14}N in the SoF might not be enough. In order to obtain the observed ratio of $^{16}\text{O}/^{18}\text{O} \lesssim 10$, H burning on the ^{12}C that accompanied the dredge-up of ^{16}O should take place, forming additional ^{14}N that will be further processed to ^{18}O .

Our attempts to simulate evolution with no driving yielded problematic results (see Section 4.2.4), which were examined in some detail using both OCTO-TIGER and FLOW-ER (see Appendix D). We conclude that with the current implementation of the ZTWD EoS,

non-driven simulations from which we can infer further physical insights are impossible. The driven simulations, on the other hand, merged before any instability started to grow, and we are confident that the merger occurred because of the unstable mass transfer (see Section 4.2.1 as well as Appendix B). We do note that the implication of the adapted EoS on the nuclear reactions and the dredge-up is unclear and should be investigated in detail in the future. We tested how a slight change in the derivation of the thermal energy (with or without the dual-energy formalism) could affect the simulation results and found only a very minor difference (Appendix E).

Finally, to accelerate OCTO-TIGER simulations we have developed a GPU-accelerated version over the last few years. We outlined our progress in Appendix F, where we showed runtimes on the GPU partition of Perlmutter using our CUDA kernels, and a noticeable GPU speedup over the CPU runs. Our aim is to expand on our results shown in this work by adding more GPU optimizations, as well as a Kokkos version, allowing us to significantly speedup future simulations and to target a wider range of supercomputers.

ACKNOWLEDGEMENTS

We thank K. J. Shen for insightful comments, and the anonymous referee for helpful remarks. This research used resources from the National Energy Research Scientific Computing Center, a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231. Portions of this research were conducted with high performance computational resources provided by the Louisiana Optical Network Infrastructure (<http://www.loni.org>). This research was supported in part by Lilly Endowment, Inc., through its support for the Indiana University Pervasive Technology Institute. This material is based upon work supported by the National Science Foundation under Grant No. CNS-0521433. This work was also supported by National Science Foundation Award 1814967. Any opinions, findings and conclusions, or recommendations expressed in this material are those of the authors, and do not necessarily reflect the views of the National Science Foundation (NSF). This work was supported in part by Shared University Research grants from IBM, Inc., to Indiana University. Octo-Tiger’s GPU development and testing was partially supported by a Swiss National Supercomputing Centre (CSCS) grant under project ID s1078. This work also required using and integrating a Python package for astronomy, yt (<http://yt-project.org>, Turk et al. 2011). JF thanks the Lorraine and Leon August Endowment to Louisiana State University for support. This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). LA-UR-24-28926.

DATA AVAILABILITY

OCTO-TIGER is available on GitHub² and was built using the following build chain.³ On Queen-Bee and BigRed OCTO-TIGER

²<https://github.com/STELLAR-GROUP/octotiger>

³<https://github.com/STELLAR-GROUP/OctoTigerBuildChain>

version Marcello et al. (2021a) was used and on NERSC's Perlmutter the pre-release state of v0.9.0 was used. We had to use different versions due to some fixes for the A100 GPUs Perlmutter (see Appendix F).

REFERENCES

- Asplund M., Gustafsson B., Lambert D. L., Rao N. K., 2000, *A&A*, 353, 287
- Bryan G. L., Norman M. L., Stone J. M., Cen R., Ostriker J. P., 1995, *Comput. Phys. Commun.*, 89, 149
- Chen H.-L., Tauris T. M., Chen X., Han Z., 2022, *ApJ*, 925, 89
- Clayton G. C., 2012, *JAASO*, 40, 539
- Clayton G. C., Geballe T. R., Herwig F., Fryer C., Asplund M., 2007, *ApJ*, 662, 1220
- Crawford C. L., Clayton G. C., Munson B., Chatzopoulos E., Frank J., 2020, *MNRAS*, 498, 2912
- Crawford C. L. et al., 2023, *MNRAS*, 521, 1674
- D'Souza M. C. R., Motl P. M., Tohline J. E., Frank J., 2006, *ApJ*, 643, 381
- Daiß G., Simberg M., Reverdell A., Biddiscombe J., Pollinger T., Kaiser H., Pflüger D., 2021, in 2021 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW). IEEE, Portland, USA, p. 377
- Daiß G. et al., 2019, in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '19. Association for Computing Machinery, New York, NY, USA
- Daiß G., Singanaboina S., Diehl P., Kaiser H., Pflüger D., 2022a, in 2022 IEEE/ACM 7th International Workshop on Extreme Scale Programming Models and Middleware (ESPM2). IEEE Computer Society, Los Alamitos, CA, USA, p. 10
- Daiß G., Diehl P., Marcello D., Kheirkhahan A., Kaiser H., Pflüger D., 2022b, in 2022 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC). IEEE, Dallas, TX, USA, p. 89
- Daiß G., Diehl P., Kaiser H., Pflüger D., 2023, in Proceedings of the 2023 International Workshop on OpenCL, IWOCCL '23. Association for Computing Machinery, New York, NY, USA
- Darwin G. H., 1879, *Proc. R. Soc. Lond. Ser. I*, 30, 1
- Diehl P. et al., 2021a, in 2021 IEEE International Conference on Cluster Computing (CLUSTER). IEEE, Portland, OR, USA, p. 204
- Diehl P. et al., 2021b, *Comput. Sci. Eng.*, 23, 73
- Diehl P., Daiß G., Huck K., Marcello D., Shiber S., Kaiser H., Pflüger D., 2023, in 2023 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW). IEEE Computer Society, Los Alamitos, CA, USA, p. 682
- Edwards H. C., Trott C. R., Sunderland D., 2014, *J. Parallel Distrib. Comput.*, 74, 3202
- Even W., Tohline J. E., 2009, *ApJS*, 184, 248
- Fryer C. L., Rockefeller G., Warren M. S., 2006, *ApJ*, 643, 292
- Han Z., 1998, *MNRAS*, 296, 1019
- Hayashi K., Kawaguchi K., Kiuchi K., Kyutoku K., Shibata M., 2021, *Phys. Rev. D*, 103, 043007
- Iben I., Jr, Tutukov A. V., 1984, *ApJS*, 54, 335
- Iben I., Jr, Tutukov A. V., 1986, *ApJ*, 311, 753
- Kaiser H. et al., 2020, *J. Open Source Softw.*, 5, 2352
- Karakas A. I., Ruiters A. J., Hampel M., 2015, *ApJ*, 809, 184
- Karambelkar V., Kasliwal M. M., Tisserand P., Clayton G. C., Crawford C. L., Anand S. G., Geballe T. R., Montiel E., 2022, *A&A*, 667, A84
- King A. R., Kolb U., 1995, *ApJ*, 439, 330
- Lai D., Rasio F. A., Shapiro S. L., 1994, *ApJ*, 423, 344
- Lauer A., Chatzopoulos E., Clayton G. C., Frank J., Marcello D. C., 2019, *MNRAS*, 488, 438
- Li Z., Chen X., Chen H.-L., Han Z., 2019, *ApJ*, 871, 148
- Lipunov V. M., Postnov K. A., 1988, *Ap&SS*, 145, 1
- Longland R., Lorén-Aguilar P., José J., García-Berro E., Althaus L. G., Isern J., 2011, *ApJ*, 737, L34
- Marcello D. et al., 2021a, *STELLAR-GROUP/octotiger: Benchmark paper*. Zenodo. Available at: <https://ui.adsabs.harvard.edu/abs/2021zndo...4432574M>
- Marcello D. C., Shiber S., De Marco O., Frank J., Clayton G. C., Motl P. M., Diehl P., Kaiser H., 2021b, *MNRAS*, 504, 5345
- Menon A., Herwig F., Denissenkov P. A., Clayton G. C., Staff J., Pignatari M., Paxton B., 2013, *ApJ*, 772, 59
- Menon A., Karakas A. I., Lugaro M., Doherty C. L., Ritter C., 2019, *MNRAS*, 482, 2320
- Montiel E. J., Clayton G. C., Marcello D. C., Lockman F. J., 2015, *AJ*, 150, 14
- Motl P. M., Tohline J. E., Frank J., 2002, *ApJS*, 138, 121
- Motl P. M., Frank J., Tohline J. E., D'Souza M. C. R., 2007, *ApJ*, 670, 1314
- Motl P. M., Frank J., Staff J., Clayton G. C., Fryer C. L., Even W., Diehl S., Tohline J. E., 2017, *ApJS*, 229, 27
- Munson B., Chatzopoulos E., Frank J., Clayton G. C., Crawford C. L., Denissenkov P. A., Herwig F., 2021, *ApJ*, 911, 103
- Paxton B., Bildsten L., Dotter A., Herwig F., Lesaffre P., Timmes F., 2011, *ApJS*, 192, 3
- Paxton B. et al., 2013, *ApJS*, 208, 4
- Paxton B. et al., 2015, *ApJS*, 220, 15
- Paxton B. et al., 2018, *ApJS*, 234, 34
- Paxton B. et al., 2019, *ApJS*, 243, 10
- Rasio F. A., 1995, *ApJ*, 444, L41
- Read J. I., Hayfield T., Agertz O., 2010, *MNRAS*, 405, 1513
- Richardson L. F., Gaunt J. A., 1927, *Phil. Trans. R. Soc. Lond. Ser. A*, 226, 299
- Ruiz-Bonilla S., Borrow J., Eke V. R., Kegerreis J. A., Massey R. J., Sandnes T. D., Teodoro L. F. A., 2022, *MNRAS*, 512, 4660
- Savitzky A., Golay M. J., 1964, *Anal. Chem.*, 36, 1627
- Schreiber M. R., Zorotovic M., Wijnen T. P. G., 2016, *MNRAS*, 455, L16
- Shen K. J., Blouin S., Breivik K., 2023, *ApJ*, 955, L33
- Staff J. E. et al., 2012, *ApJ*, 757, 76
- Staff J. E. et al., 2018, *ApJ*, 862, 74
- Timmes F. X., Swesty F. D., 2000, *ApJS*, 126, 501
- Tisserand P. et al., 2022, *A&A*, 667, A83
- Turk M. J., Smith B. D., Oishi J. S., Skory S., Skillman S. W., Abel T., Norman M. L., 2011, *ApJ*, 192, 9
- Tutukov A. V., Yungelson L. R., 1994, *MNRAS*, 268, 871
- Webbink R. F., 1984, *ApJ*, 277, 355

APPENDIX A: ZERO-TEMPERATURE WHITE DWARF EQUATION OF STATE

When using the zero-temperature white dwarf (ZTWD) equation of state (EoS) in OCTO-TIGER, the fluid is modelled using a combination of a zero-temperature Fermi fluid and an ideal gas. The total pressure, P , is the sum of pressure from the zero-temperature fluid, P_{deg} , and the ideal gas internal (thermal) pressure, P_{th} . The gas energy, the quantity OCTO-TIGER evolves, is the sum of internal (thermal), degenerate, and bulk kinetic energy densities.

$$E = E_{\text{th}} + E_{\text{deg}} + E_{\text{kin}}. \quad (\text{A1})$$

The energy density of the degenerate electron gas, E_{deg} , is given by the relation,

$$E_{\text{deg}} = H_{\text{deg}} \rho - P_{\text{deg}}. \quad (\text{A2})$$

where

$$H_{\text{deg}} = (8A/B) \left[(x^2 + 1)^{1/2} - 1 \right] \quad (\text{A3})$$

is the specific enthalpy of the degenerate electron gas,

$$P_{\text{deg}} = A \left[x (2x^2 - 3) (x^2 - 1)^{1/2} - 3 \sinh^{-1} x \right], \quad (\text{A4})$$

and

$$x \equiv (\rho/B)^{1/3}. \quad (\text{A5})$$

The constants A and B are

$$\begin{aligned} A &= \pi m_e^4 c^5 / (3h^3) = 6.00228 \times 10^{22} \text{ dynes cm}^{-2}, \\ B &= \mu_e \times 8\pi m_p m_e^3 c^3 / (3h^3) = \mu_e \times 9.81011 \times 10^5 \text{ g cm}^{-3}, \end{aligned} \quad (\text{A6})$$

where μ_e is the average ratio of nucleons to electrons. We assume that throughout each white dwarf (and basically throughout the whole domain) $\mu_e = 2$, hence, $B = 1.96202 \times 10^6 \text{ g cm}^{-3}$.

For numerical reasons, at low density regions ($x < 0.001$), we approximate E_{deg} by a Taylor expansion

$$E_{\text{deg}} (x < 0.001) = 2.4 \times Ax^5. \quad (\text{A7})$$

As described in Marcello et al. (2021b), OCTO-TIGER evolves a second variable for the energy, the ‘entropy tracer’, τ (Motl, Tohline & Frank 2002). The inclusion of this additional variable allows for the proper evolution of shocks while still retaining a precise calculation of the thermal energy in regions where the kinetic energy dominates over the non-degenerate energy components. The thermal energy density therefore is computed according to

$$E_{\text{th}} = \begin{cases} E - E_{\text{kin}} - E_{\text{deg}}, & \text{if } E - E_{\text{kin}} - E_{\text{deg}} \geq \epsilon_1 E \\ \tau^\gamma, & \text{otherwise} \end{cases}, \quad (\text{A8})$$

where $\gamma = \frac{5}{3}$ is the ratio of specific heats. In addition, the entropy tracer is reset using $\tau = (E_{\text{th}})^{\frac{1}{\gamma}}$, in computational cells which satisfy

$$E - E_{\text{kin}} - E_{\text{deg}} > \epsilon_2 E \quad (\text{A9})$$

for at least one of the adjacent cells or the cell itself. Otherwise it is left alone. Note the additional degenerate energy terms in A8 and A9 compared to what is described in Marcello et al. (2021b). Throughout this paper, we use the default values $\epsilon_1 = 0.001$ and $\epsilon_2 = 0.1$ as was chosen by Bryan et al. (1995).

The combination of using both a ZTWD EoS and a dual energy formalism implies that in the WD interiors, where the degenerate energy dominates, the entropy tracer is not directly updated from the gas energy and is only affected by diffusion. Initially, the entropy tracer is set to a very low value inside the stars, which corresponds to a very low, practically zero, temperature. However, as the simulation evolves in time the entropy tracer begins to increase inside the stars, which effectively slightly heats the stars. To test whether this heating is significant, we have run a simulation where we disable the dual energy formalism and found that this has only a very little effect on the results (see Appendix E).

Finally, the temperature is calculated as

$$T = \frac{(\gamma - 1) E_{\text{th}}}{n K_B}, \quad (\text{A10})$$

where

$$n = \sum_i \frac{\rho_i}{\mu_i m_H}, \quad (\text{A11})$$

μ_i is the molecular weight of the i th specie, m_H is the mass of an hydrogen atom, and K_B is the Boltzmann constant. The total pressure is the sum of the degenerate and thermal pressure components

$$P = P_{\text{deg}} + (\gamma - 1) E_{\text{th}}, \quad (\text{A12})$$

APPENDIX B: VERIFICATION OF THE EARLY MASS TRANSFER WITH ANALYTICAL CONSIDERATIONS

In what follows we analyse the results of our numerical simulations focusing on the results of the highest resolution L13 run, with two goals in mind: a check of the accuracy of the simulations, and a better understanding of the evolution of the binary system. We start from an approximate analytical expression for the total angular momentum of the binary as the sum of the two-body orbital angular momentum for point-mass components plus the contribution from the spins, where I_1 and I_2 are the moments of inertia of the components with respect to their respective centers of mass:

$$J = M_1 M_2 \sqrt{\frac{Ga}{M}} (1 - e^2)^{\frac{1}{2}} + I_1 \Omega_1 + I_2 \Omega_2. \quad (\text{B1})$$

This expression is approximate in that we are using the simple Keplerian two-body orbital frequency for point masses, valid if both components remain spherically symmetrical. This approximation gets gradually worse as the binary separation is reduced and tidal distortions increase. We construct from the above the following logarithmic derivatives:

$$\begin{aligned} \frac{\dot{J}}{J_{\text{orb}}} &= \frac{\dot{M}_1}{M_1} + \frac{\dot{M}_2}{M_2} + \frac{1}{2} \frac{\dot{a}}{a} - \frac{1}{2} \frac{\dot{M}}{M} - \frac{e\dot{e}}{(1 - e^2)} \\ &\quad + \frac{\dot{I}_1 \Omega_1 + I_1 \dot{\Omega}_1 + \dot{I}_2 \Omega_2 + \Omega_2 \dot{\Omega}_2}{J_{\text{orb}}}, \end{aligned} \quad (\text{B2})$$

where the part of the right-hand-side that is on the first line is the $\dot{J}_{\text{orb}}/J_{\text{orb}}$, while the term that is on the second line is $\dot{J}_{\text{spin}}/J_{\text{orb}}$:

$$\frac{\dot{J}_{\text{spin}}}{J_{\text{orb}}} = \frac{J_1}{J_{\text{orb}}} \left(\frac{\dot{I}_1}{I_1} + \frac{\dot{\Omega}_1}{\Omega_1} \right) + \frac{J_2}{J_{\text{orb}}} \left(\frac{\dot{I}_2}{I_2} + \frac{\dot{\Omega}_2}{\Omega_2} \right). \quad (\text{B3})$$

Finally, we set an explicit expression for the left-hand-side of equation (B2):

$$\dot{J} = \dot{J}_{\text{sys}} + \dot{M} j_{L2}, \quad (\text{B4})$$

where \dot{J}_{sys} represents systemic angular momentum losses, either numerical (see below) or real, such as gravitational radiation and magnetic braking, and j_{L2} is the effective specific angular momentum of gas at L2. In our simulations the initial driving is set to $\dot{J}_{\text{sys}} = -J_{\text{orb}} \times 0.01/P_0$ for a time $0 \leq t \leq T_{\text{drv}}$, and zero for $t > T_{\text{drv}}$. Mass-loss from the binary occurs mostly through L2, but also through L3, and gravitational coupling between the binary and the mass outflow will affect the average specific angular momentum of the outflow. Thus j_{L2} should




Figure B1. The graph shows the instantaneous differences between logarithmic derivatives on the left- and right-hand sides of the equation (B2), which should be zero at all times. Rolling averages of this value are also plotted. The red vertical lines are each placed one orbital period apart, and the blue vertical line (at $t \sim 260$ seconds) indicates the time at which driving ceased.

be regarded as an effective specific angular momentum of the outflow.

While the above equations are in principle valid for the instantaneous values of the masses and orbital parameters, in practice the determination of these quantities and their derivatives from the numerical simulation, requires averaging over data points covering a time $\sim P_0$, where P_0 is the initial orbital period. Thus these equations should be considered effectively orbit-averaged equations. To construct the various terms in equation (B2), we obtain at every instant t the best-fitting elliptical orbit, based on the data points for the interval $(t - P_0, t)$. This fit yields the orbit-averaged values for the separation a , the eccentricity e , and their derivatives, and OCTO-TIGER provides the rest of the quantities.

Fig. B1 shows the instantaneous (blue) difference between the logarithmic derivatives on the left- and right-hand sides of equation (B2) based on the results of simulation L13. The orange and green curves represent the above differences averaged over 20 and 2000 diagnostic samples, where a sample is approximately five time-steps. Perhaps the most remarkable message conveyed by this figure is how well the orbit-averaged equations describe the numerical evolution of the binary which is free of the simplifying assumptions of the analytic model. Significant deviations from the simple model only occur during the initial driving and the final dynamical merger. Tidal distortions of both components, mass-loss through the outer Lagrange points, and ultimately the inability to define unequivocally what constitutes the donor during its disruption, ultimately make the simple model inapplicable. The fact that the numerical simulation matches the expectations from the simple model precisely when the eccentricity is small ($e \approx 0.01$, see top left panel in Fig. 6) and the components are nearly spherical strengthens our confidence in the validity of our results.

APPENDIX C: BINARY EVOLUTION COMPARISON BETWEEN FLOW-ER AND OCTO-TIGER

We find that the FLOW-ER simulations are very comparable to the high resolution OCTO-TIGER simulations with respect to resolution and to the amount of driving. The L12 simulation has a similar number of cells across the donor to simulations FL-1 and FL-2, while the L13 simulation has a similar total number of cells in the domain, but higher spatial resolution of the donor (Table 1). The only difference between FL-1 and FL-2 is that the former is driven for one orbit, while the latter is driven for two, which is comparable to L12, which is driven for 1.3 orbits, and to L13 which is driven for 2.3 orbits.

In Fig. C1 we show a set of diagnostic quantities as a function of time for the four simulations, similar to what we show in Fig. 11. The merging times compare as one would expect: the simulations driven for ~ 2 orbit merge at 10 and 17 orbits, for the L13 and FL-1, respectively, while those driven for ~ 1 orbits merge at 37 and 39 orbits, for the L12 and FL-2, respectively.

The only anomaly in behaviour is for L13, which exhibits a decrease in the central density of the accretor before the merger, not exhibited by any other simulation in this comparison. Most noticeable is the extremely consistent mass transfer rate onto the accretor for all four simulations.

Below are some links to the movies showing density slices of the FLOW-ER simulations:

- [Equatorial Slice of Q0.7 driven for 1 orbit](#)
- [Meridional Slice of Q0.7 driven for 1 orbit](#)
- [Equatorial Slice of Q0.7 driven for 2 orbits](#)
- [Meridional Slice of Q0.7 driven for 2 orbits](#)




Figure C1. Comparison between L12 and L13 and the two FLOW-ER simulations. Orbital separation and orbital angular momentum (first row), mass ratio, and mass transfer rate (second row), maximum accretor density and donor angular momenta (third row).

APPENDIX D: NON-DRIVEN SIMULATIONS – HITTING THE NUMERICAL BARRIER

When trying to simulate a non-driven ‘more realistic’ evolution one should expect a long (in wall-clock time) and expensive (in computer resources) simulation. Moreover, on hundreds (or thousands) of orbits numerical errors that accumulate can come into effect. For example, instead of the actual mass transfer (which can be very slow), deviations from angular momentum conservation can start dictating the evolution. We wanted to take advantage of the excellent angular momentum conservation in OCTO-TIGER (see Fig. 2, for example), and the fact that OCTO-TIGER can scale very well to a high number of cores to try and run non-driven simulations in a reasonable wall clock time. Unfortunately, as we will describe below, we have hit another numerical barrier, which has never been reached before. This time, the issues stem from the specific usage of the ZTWD EoS.

In Table D1 we present the non-driven simulations L11ND, L12ND that are otherwise identical to the previously discussed L11, L12. Diagnostic plots for these four simulations are shown in Fig. D1.


At first glance, the evolution of the orbital angular momentum agrees very well and again shows remarkable resemblance, slowly declining through most of the evolution but then rapidly dropping just before the merger for simulations L11, L11ND, and L12 (we stopped L12ND before it merged). Also, with no surprise, the non-driven L11ND merged much later than the driven L11.

However, in simulation L11ND, the primary (the more massive star that should be the accretor) is donating mass to the secondary (less massive star) and the mass ratio grows above 0.7 before rapidly decreasing right before the merger. Likewise, the mass transfer rate between donor and accretor is positive for driven simulations with the mass accretion rate in L11ND, L12ND becoming negative.

Table D1. Our non-driven simulations.

Source	Code	M_{tot} (M_{\odot})	EoS	T_{drv} (P_0)	T_{merge} (P_0)	Resolution; Δx_{min} (R_{\odot}) (cells/particles)	Donor res (radial cells)	a_0 (10^9 cm)	P_0 (s)	Hybrid
L11ND	Octo (Bd)	0.9	ZTWD	0	51.4	2.5 M; 5.73×10^{-4}	73	3.413	114.0	Yes
L12ND	Octo (Bd)	0.9	ZTWD	0	–	5.3 M; 2.86×10^{-4}	145	3.403	113.6	Yes

Note. Notations are as in Table 1.


Figure D1. Comparison between the non-driven simulations L11ND, and L12ND and the driven simulations L11, and L12.

Furthermore, L11ND and L12ND are clearly anomalous with the central density of the primary declining (and the potential energy increasing) throughout the simulation. The accretor expands in these non-driven simulations, overflows its Roche lobe, and transfers mass to the secondary star, which ultimately triggers the merger. The expansion itself is caused by an instability, convection-like, that is forming inside the accretor. We stopped L12ND the moment we witnessed the expansion to save computer time.

To isolate this problem, we have tested whether we can reproduce this expansion in a single star simulation. In this test, we created a hydrostatic model of a white dwarf using the ZTWD EoS. We have run two resolutions, which correspond to the resolutions of the L10ND and L11 binary simulations, with respect to the number of cells across the star's diameter. We fixed the star mass to the mass of the accretor and thus the star structure closely follows the initial accretor structure in our binary simulations. Like the stationary star test in Marcello et al. (2021b), the grid size was twice the size of the




Figure D2. Internal convection-like instabilities in a single star simulation. Density (left), thermal pressure divided by total pressure (centre), and temperature (right) slices along the equatorial plane are plotted. We show equi-potential contours and velocity arrows on the left and centre panels, while on the right panel, we show density contours. Hot bubbles that form inside the stars buoy outward, disrupting the stability of the star.




Figure D3. Stability tests of a single star with a ZTWD EoS. From left to right: The central density, the maximal radius of a density contour equal to 10^3 g cm^{-3} , and the relative energy error (see the text) as a function of time. On the higher resolution, using the dual-energy formalism (DEF), the star starts to expand due to internal, convection-like instabilities, at about 40 initial periods P_0 of the $q = 0.7$ binary (solid thick line). Disabling the DEF prolongs the stability of the star and postpones the sudden decrease in the central density to nearly 80 P_0 (or ~ 2600 dynamical times of the single star; dashed thick line).

star. On low resolution, the star remains stationary on thousands of dynamical times (the initial orbital period, P_0 , shown in the plots is 114 sec and is equal to 33 dynamical times). This gives us further confidence in our L10ND simulation, where such an expansion was not observed. However, on the L11 equivalent resolution (64 cells across), the same instabilities that developed in the binary simulations L11ND, and L12ND are observed.

We illustrate these internal convection-like instabilities in Fig. D2. We plot density (left), thermal pressure divided by total pressure (center), and temperature (right) at the time they begin to form, at approximately $32P_0$. Each panel shows a slice through the equatorial plane from the single star simulation where we resolve 64 cells across the star’s diameter and where the thermal energy is computed according to the dual-energy formalism (DEF; equation A8). On the most left and central panels, we also plot equi-potential contours and velocity arrows, while on the right panel, we plot density contours. Bubbles of hotter temperature than their surroundings are clearly seen in the interiors of the star at radii of $\sim 0.2 - 0.5 \times 10^9 \text{ cm}$. In these bubbles, the thermal pressure is also higher compared to the thermal pressure of their surroundings (the *total pressure* is not that different though). Those bubbles tend to rise outward in a convection like motion, disrupting the structure of the star. As a consequence, instabilities are formed along the star surface, which results in a complex flow. We found out that this behaviour precedes the sudden rapid expansion of the star.

In Fig. D3 we plot the central density (left), the maximal radial distance to the isopycnic surface $\rho = 10^3 \text{ g cm}^{-3}$ (middle), and the deviation from conservation of energy (right): $\Delta E = (E(t) - E(0))/E(0)$ as a function of time, for several runs using two spatial resolutions with and without DEF. Here $E(t) = \int (E_{\text{deg}} + E_{\text{th}} + E_{\text{kin}} + \Phi) dV$ over the entire grid. In contrast with Fig. 2, here we do not include outflows of energy in the calculation. Along with the effects of resolution and the usage of the DEF, we ran a case where only degenerate pressure is included (the thermal pressure is zero and the total pressure equal to the degenerate pressure).

The FLOW-ER method, (without the DEF; i.e. equation E1), maintains the star structure for longer, postponing the instabilities to a much later phase in the evolution. However, eventually, the star expands and the total energy deviation grows to 30 percent. The instabilities are completely avoided in the simulation that includes degenerate pressure only and the star remains intact throughout the evolution. This further confirms the role of the convection-like instabilities in the expansion and the decrease in the central density. It also pinpoints the problem of the way we add the contribution of the thermal pressure to the degenerate pressure. A similar stationary test with an ideal gas pressure only does not show the same instability (see the stationary star test in Marcello et al. 2021b). Moreover, an equivalent run performed with FLOW-ER including both degenerate and thermal pressure contributions, experiences the same convection-

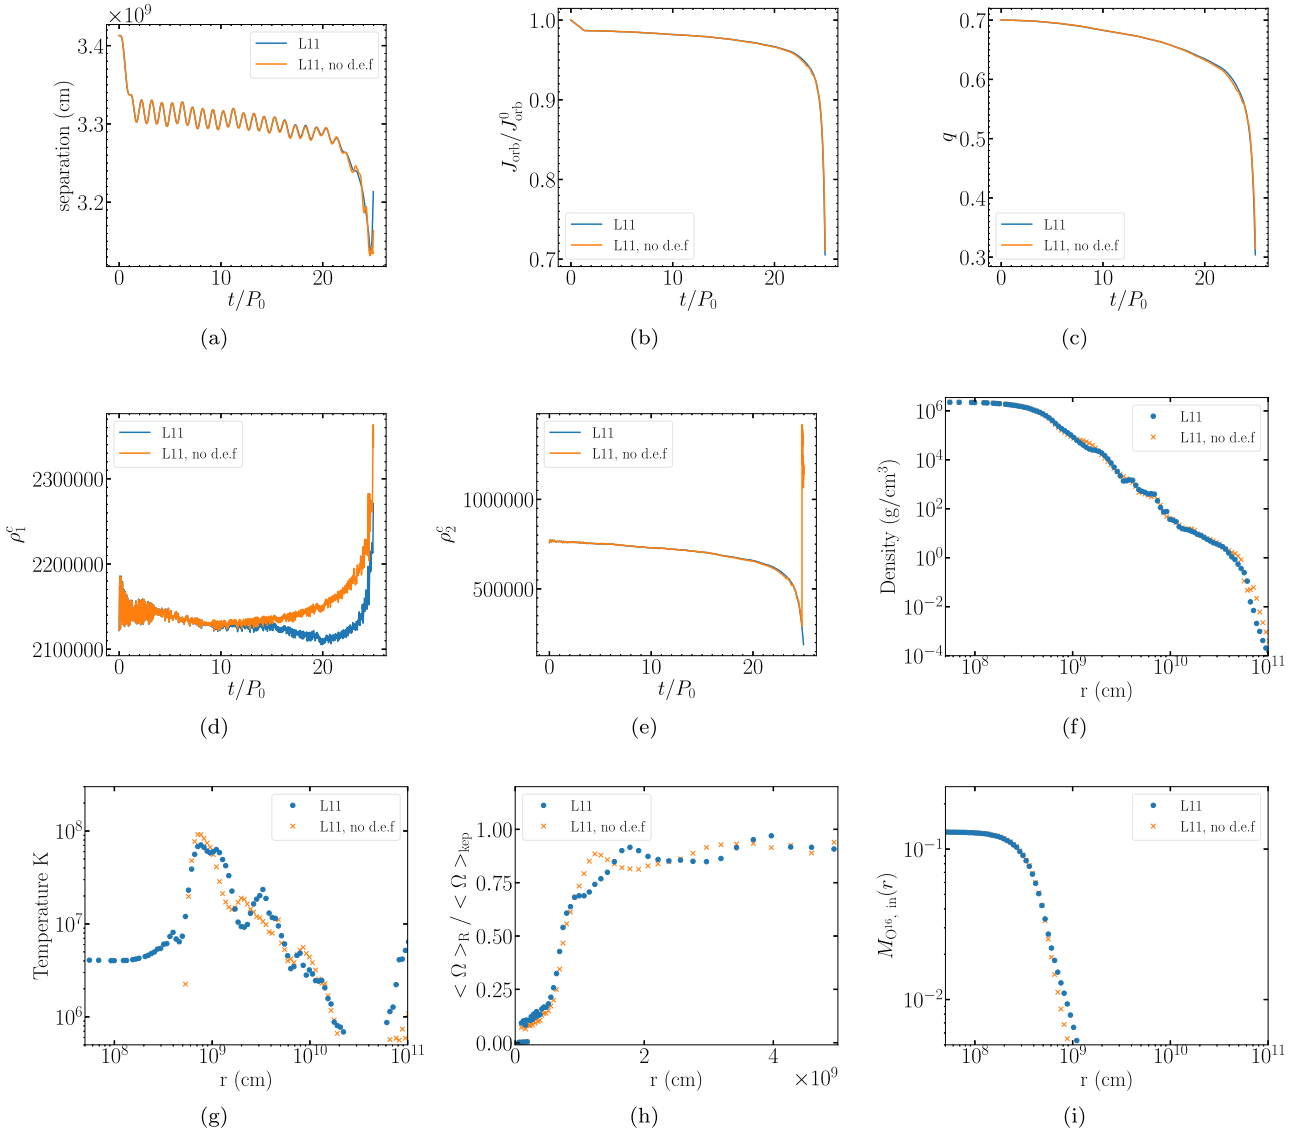


Figure E1. Comparing two driven simulations with the same resolution but with different treatments for the derivation of the internal (thermal) energy. L11 is the simulation shown throughout the paper in which the thermal energy is calculated according to equation (A8). We use this treatment in all the simulations shown in this paper. In the other simulation, we used instead equation (E1), similar to what is being performed in FLOW-ER. Only very minor differences exist which indicate that the specific temperature prescription does not affect much the simulations.

like instability, although much later, and results in the expansion of the star and the decrease of the central density. It is possible that using a Helmholtz EoS, which takes into account the varying degeneracy level of the electrons, rather than the ZTWD which only assumes a full degenerate electron pressure with thermal ions could resolve this problem. However, that is outside the scope of this paper. In addition, it is important to note that previous grid-based simulations using the ZTWD (like in Staff et al. 2012 or Staff et al. 2018), have never been run long enough to reveal this phenomenon.

APPENDIX E: THE EFFECT OF DISABLING THE DUAL ENERGY FORMALISM

To check the effect of using the dual-energy formalism (DEF) on the short-driven simulations we performed a simulation where the DEF is disabled. The thermal energy in a cell is calculated then by

subtracting the kinetic plus degenerate energies from the gas energy while forcing it to be non-negative,

$$E_{\text{th}} = \max(E - E_{\text{deg}} - E_{\text{kin}}, 0). \quad (\text{E1})$$

In this simpler treatment, the entropy tracer is disregarded and never used to calculate the thermal energy. This approach is used in all the FLOW-ER simulations shown in this paper. We have run an equivalent of the L11 driven simulation just with the FLOW-ER approach (equation E1) and we demonstrate in Fig. E1 that these simulations appear to be very similar.

First, we note that the orbital properties of the system evolve with time almost identically in both of the simulations. The orbital separation (a), angular momentum (b), and the mass ratio (hence the mass transfer rate) (c) evolve as described in 4.1 and behave as expected, with their respective curves falling on top of each other most of the time. The donor central density (e) also behaves similarly, indicating that the donor star structure is not affected by the DEF




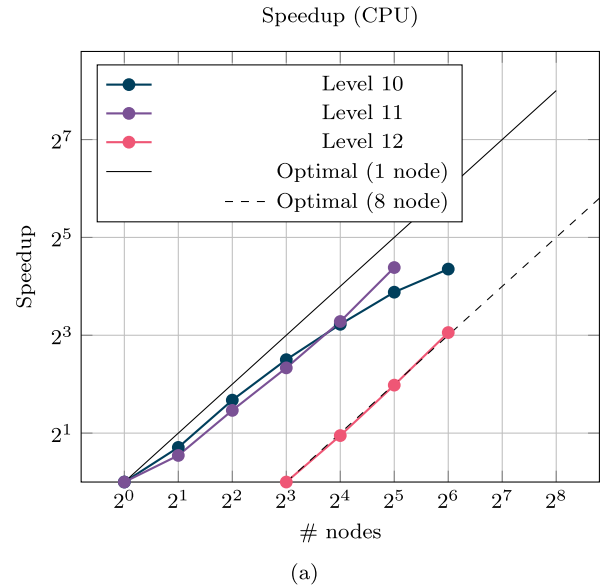
Figure F1. Scaling on NERSC's Perlmutter for simulations for runs using CPU only and using CPU and GPU. All runs started on a single node, and we increased the number of nodes until we did not do sufficient work and the scaling flattened out. However, level 12 only fitted on eight nodes. The configuration on Perlmutter was as follows: We used four localities with 16 cores per locality for the CPU-only run. One NVIDIA[®] A100 was added to each locality for the combined CPU and GPU runs.

treatment being used. The most notable difference is in the accretor's central density (d), which decreases more in the simulation where the DEF is used. This results in a slightly less dense core at the time of the merger. However, the averaged density structure of the merger five orbits after the merger (f), almost converges, with a small bump at $\sim 10^9$ cm as the only difference. The temperature and rotation profiles, (g) and (h), are also very close besides small variations. Finally, the accretor expands slightly further prior to the merger in the simulation with the DEF. Although this results in a little more diffusion of oxygen-16 out (i) and displaces the first peak in angular velocity slightly outward (h), overall, the effect is marginal.

APPENDIX F: SCALING AND CODE PERFORMANCE USING NVIDIA GPUS

OCTO-TIGER's support for GPU acceleration within the C++ standard library for parallelism and concurrency (HPX) (Kaiser et al. 2020) was recently extended. Initially, GPU support was added (via CUDA) when OCTO-TIGER's gravity module was ported to the GPU and tested on CSCS's Piz Daint (Daiß et al. 2019). Recently, OCTO-TIGER's new hydro module using a higher order reconstruction was similarly ported and tested on ORNL's Summit in Diehl et al. (2021a). For these simulations, we used HPX's native CUDA implementation.

Now, to investigate the speed-up brought by the NVIDIA A100 GPUs on NERSC's Perlmutter, we executed our simulations shortly before the merger. Thus, we have to have a very refined irregular grid. Each time, the code was run for 25 time-steps for level 10 and level 11. For level 12, the runs on a single node without GPUs got very time-consuming, and we reduced the time steps from 25 to five time-steps. Table F1 shows the number of cells and memory usage of each of the simulations. Fig. F1 shows the scaling from close to the



Speedup (CPU + GPU)

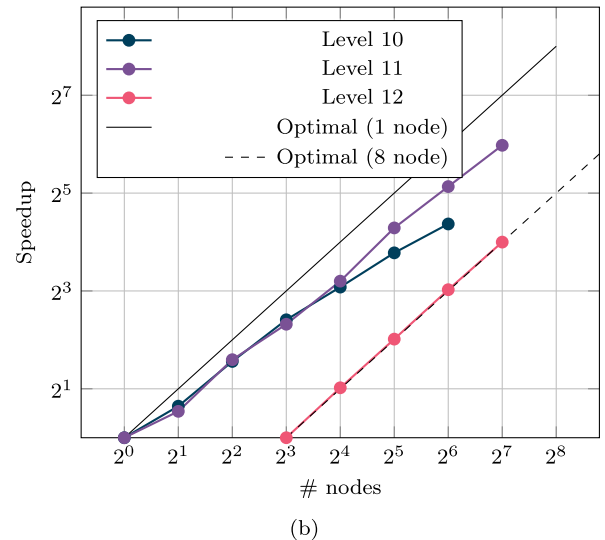


Figure F2. Speedup for the runs from close to the merger on NERSC's Perlmutter from one to 128 nodes. In (a) only the CPUs were used, and in (b) the NVIDIA[®] A100 GPUs were added. Note that in the Phase 1, the maximal number of nodes was restricted to 128.

Table F1. Number of cells and memory usage for simulations L10ND, L11ND, and L12 at the starting point of our scaling runs. We use a time earlier to the merger, when massive regridding takes place, as our starting point.

Level	Number of cells	Memory
L10ND	3.8M	11 GB
L11ND	40.2M	113 GB
L12	257.3M	724 GB

Table F2. Toolchain and OCTO-TIGER’s dependencies on Perlmutter.

gcc	9.3.0	hwloc	1.11.12
cray-mpich	8.1.11	boost	1.77.0
CUDA™	11.4.0	jemalloc	5.1.0
hpx	1.7.1	silos	4.10.2
hdf5	1.8.12	cppuddle	d32e50b

Table F3. NERSC’s Perlmutter configuration.

CPU	AMD® EPYC 7713 64-Core Processor
GPU	4 × NVIDIA® A100-PCIE-40GB
GPU driver	450.162
Linux kernel	5.3.18

merger for 25 time-steps. Fig. F2 shows the corresponding speedup. We used four A100 GPUs per node and for each GPU, we used one MPI rank with 16 cores assigned. So in total, we used all 64 cores per node. Tables F2 and F3 show the software version and Perlmutter’s hardware, respectively. The dependency cppuddle⁴ is available on GitHub.

With the ongoing heterogeneity of acceleration cards in the latest supercomputers, e.g. AMD GPUs in ORNL’s Frontier, NVIDIA GPUs in NERSC’s Perlmutter, and hopefully soon Intel GPUs in Aurora, we prepared OCTO-TIGER to support Kokkos (Edwards, Trott & Sunderland 2014) to target various GPU vendors. The Kokkos

integration (Daiß et al. 2021) was tested on Riken’s Supercomputer Fugaku (Diehl et al. 2023) and Stony Brook’s Ookami using Arm A64FX CPUs with SVE vectorization (Daiß et al. 2022a). Recently, SYCL support was added to Octo-Tiger to target Intel GPUs (Daiß et al. 2023). Additionally, support for dynamic GPU work aggregation was added in Daiß et al. (2022b) to improve the GPU performance further. These improvements and additions prepare OCTO-TIGER for a diverse set of acceleration cards. Significantly, the improved speedup using GPUs is promising since we have shown in this paper that extensive simulations are needed for future research.

Disclaimer

The results on NERSC’s Perlmutter were conducted in phase 1, and as such these results should not reflect or imply that they are the final results of the system. Numerous upgrades will be taking place for Phase 2 that will substantially change the final size and network capabilities of Perlmutter.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.

⁴<https://github.com/SC-SGS/Cppuddle>