

Lawrence Berkeley National Laboratory

LBL Publications

Title

Dynamics of Laterally Propagating Flames in X-Ray Bursts. II. Realistic Burning and Rotation

Permalink

<https://escholarship.org/uc/item/5r68v4x0>

Journal

The Astrophysical Journal, 912(1)

ISSN

0004-637X

Authors

Harpole, A
Ford, NM
Eiden, K
[et al.](#)

Publication Date

2021-05-01

DOI

10.3847/1538-4357/abee87

Peer reviewed

Dynamics of Laterally Propagating Flames in X-ray Bursts. II. Realistic Burning & Rotation

A. HARPOLE,¹ N. M. FORD,² K. EIDEN,^{3,1} M. ZINGALE,¹ D. E. WILLCOX,² Y. CAVECCHI,⁴ AND M. P. KATZ⁵

¹*Dept. of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794-3800*

²*Lawrence Berkeley National Laboratory, Berkeley, CA*

³*Dept. of Astronomy, University of California, Berkeley, CA 94720-3411*

⁴*Universidad Nacional Autónoma de México, Instituto de Astronomía, Ciudad Universitaria, CDMX 04510, Mexico*

⁵*NVIDIA Corporation*

ABSTRACT

We continue to investigate two-dimensional laterally propagating flames in type I X-ray bursts using fully compressible hydrodynamics simulations. In the current study we relax previous approximations where we artificially boosted the flames. We now use more physically realistic reaction rates, thermal conductivities, and rotation rates, exploring the effects of neutron star rotation rate and thermal structure on the flame. We find that at lower rotation rates the flame becomes harder to ignite, whereas at higher rotation rates the nuclear burning is enhanced by increased confinement from the Coriolis force and the flame propagates steadily. At higher crustal temperatures, the flame moves more quickly and accelerates as it propagates through the atmosphere. If the temperature is too high, instead of a flame propagating across the surface the entire atmosphere burns uniformly. Our findings could have implications for the relationship between observed burst rise times and neutron star rotation and accretion rates. All of the software used for these simulations is freely available.

Keywords: X-ray bursts (1814), Nucleosynthesis (1131), Hydrodynamical simulations (767), Hydrodynamics (1963), Neutron stars (1108), Open source software (1866), Computational methods (1965)

1. INTRODUCTION

Considerable evidence suggests that ignition in an X-ray burst (XRB) starts in a localized region and then spreads across the surface of the neutron star (Bhattacharyya & Strohmayer 2007; Chakraborty & Bhattacharyya 2014). We continue our study of flame spreading through fully compressible hydrodynamics simulations of the flame. Building on our previous paper (Eiden et al. 2020), we relax the approximations we made previously (artificially boosting the speed of the flame in order to reduce the computational cost) and explore how the flame properties depend on rotation rate and the thermal structure of the neutron star. This new set of realistic simulations is possible because of the work done to offload our simulation code, Castro (Almgren et al. 2020), to GPUs, where it runs significantly faster.

We investigate the effect of rotation rate on the flame. With the exception of IGR J17480-2446 (Altamirano et al. 2010, spinning at 11 Hz), most observations of XRBs which come from sources with known rotation rates have rotation rates of 200 – 600 Hz (Bilous & Watts 2019; Galloway et al. 2020). There are a number of factors that could explain this lack of observations below 200 Hz. It could be that there is some physical process which inhibits the flame ignition and/or spread at lower rotation rates. It could be that bursts at lower rotation rates are smaller in amplitude and therefore more difficult to detect. It could be that it does not have anything to do with the flame at all, but that neutron stars in accreting low mass X-ray binaries rarely have rotation rates below 200 Hz.

Previous studies have found that rotation can have a significant effect on the flame’s propagation. As the rotation rate increases, the Coriolis force whips the spreading flame up into a hurricane-like structure (Spitkovsky et al. 2002; Cavecchi et al. 2013). The stronger Coriolis force leads to greater confinement of the hot accreted matter, leading to easier ignition of the flame (Cavecchi et al. 2015).

The temperature structure of the accreted fuel layer can also affect the flame propagation. Timmes (2000) showed that laminar helium flames have higher speeds when moving into hotter upstream fuel. It has been suggested that crustal heating may be stronger at lower accretion rates and weaker at higher accretion rates, due to the effect of neutrino losses (Cumming et al. 2006; Johnston et al. 2019). On the other hand, at very high accretion rates the atmosphere is so heated that it simmers in place rather than forming a propagating flame (Fujimoto et al. 1981; Bildsten 1998; Keek et al. 2009). A shallow heating mechanism of as yet unknown origin has been found necessary to reproduce observed properties of XRBs in 1D simulations (Deibel et al. 2015; Turlione et al. 2015; Keek & Heger 2017). In our models, we keep the crust at a constant temperature, so by increasing this temperature we can effectively increase the crustal heating, shallow heating and/or mimic the effects of accretion-induced heating.

In the following sections, we conduct a series of simulations at various rotation rates and crustal temperatures to investigate their effects on the flame. We find that at lower rotation rates, the flame itself becomes harder to ignite. At higher rotation rates, nuclear burning is enhanced and the flame propagates steadily. At higher crustal temperatures, burning is greatly enhanced and the flame accelerates as it propagates. We discuss the implications that this may have for burst physics and observations.

2. NUMERICAL APPROACH

We use the *Castro* hydrodynamics code (Almgren et al. 2010; Almgren et al. 2020) and the simulation framework introduced in Eiden et al. (2020). The current simulations are all performed in a two-dimensional axisymmetric geometry. For these axisymmetric simulations, we add an additional geometric source term from Bernard-Champmartin et al. (2012) that captures the effects of the divergence of the flux operating on the azimuthal unit vector. This term is a small correction, but was missing from our previous simulations. The simulation framework initializes a fuel layer in hydrostatic equilibrium, laterally blending a hot model on the left side of the domain (the coordinate origin) and a cool model on the right. The initial temperature gradient between the hot and cool fluids drives a laterally propagating flame through the cool fuel. In our original set of calculations (Eiden et al. 2020), in order to make the simulations computationally feasible we artificially boosted the flame speed by adjusting the conductivity and reaction rate to produce a flame moving 5–10 \times faster than the nominal laminar flame speed. We also used high rotation rates (≥ 1000 Hz) to reduce the lateral lengthscale at which the Coriolis force balances the lateral flame spreading in order to reduce the size of the simulation domain. The port of *Castro* to GPUs (Katz et al. 2020) significantly improved its overall performance, enabling us to run these new simulations without the previous approximations while continuing to resolve the burning front. For these simulations, we no longer boost the flame speed—the true




Figure 1. Initial temperature structure (left panel) and density structure (right panel) as a function of height in the “hot” state.

conductivities (taken from Timmes 2000) and reaction rates are used. We are also able to use slower, more physically realistic rotation rates. We continue to use a 13-isotope α -chain to describe the helium burning.

The initial model is set up in the same fashion as described in Eiden et al. (2020). In particular, we create a “hot” and “cool” hydrostatic model representing the ash and fuel states and blend the two models laterally to create a hot region near the origin of the coordinates and a smooth transition to the cooler region at larger radii. This hot region extends out to $r_{\text{pert}} = 4.096 \times 10^4 \text{ cm}$. The cool initial model is characterized by three temperatures: T_{star} is the isothermal temperature of the underlying neutron star, T_{hi} is the temperature at the base of the fuel layer, and T_{lo} is the minimum temperature of the atmosphere. The atmosphere structure is isentropic as it falls from T_{hi} down to T_{lo} , and so is marginally unstable to convection. For the hot model, we replace T_{hi} with $T_{\text{hi}} + \delta T$. In the calculations presented here, we explore the structure of the initial models by varying these parameters. All models have the same peak temperature in the hot model, $T_{\text{hi}} + \delta T$.

For the current simulations, we explore a variety of initial rotation rate and temperature conditions for the flame. The main parameters describing the models and the names by which we shall refer to them in this paper are provided in Table 1. Figure 1 shows the temperature and density structure for our hot models and Figure 2 shows the temperature and density structure for the cool models.

3. SIMULATIONS AND RESULTS

We present six simulations in total, summarized in Table 1. These simulations encompass three different rotation rates: 250 Hz, 500 Hz, and 1000 Hz, and for the 500 Hz run, four different temperature profiles. In the following subsections, we look at how the flame properties depend on the model parameters. All simulations are run in a domain of $1.8432 \times 10^5 \text{ cm} \times 3.072 \times 10^4 \text{ cm}$ with a coarse grid of 1152×192 zones and two levels of refinement (the first level refining the resolution by factor of four, and the second by a factor of two again). This gives a fine-grid resolution of 20 cm. In these simulations, refinement is carried out in all zones within the atmosphere with density $\rho > 2.5 \times 10^4 \text{ g cm}^{-3}$. We use an axisymmetric coordinate system, with the horizontal r -direction pointing along the surface of the star and the vertical z -direction pointing perpendicular to the surface.




Figure 2. Initial temperature structure (left panel) and density structure (right panel) as a function of height in the “cool” state.

Table 1. Rotation rate and temperature properties of the simulations. In the left-hand column we list the names we shall use to refer to each simulation throughout this paper.

run	Rotation Rate (Hz)	δT (K)	T_{hi} (K)	T_{star} (K)	T_{lo} (K)
F1000	1000	1.2×10^9	2×10^8	10^8	8×10^6
F500	500	1.2×10^9	2×10^8	10^8	8×10^6
F500_2E8	500	1.2×10^9	2×10^8	2×10^8	8×10^6
F500_3E8	500	1.1×10^9	3×10^8	3×10^8	8×10^6
F500_4E8	500	10^9	4×10^8	4×10^8	8×10^6
F250	250	1.2×10^9	2×10^8	10^8	8×10^6

For some of our analysis, we would like to have a means of estimating the temperature (T) and nuclear energy generation rate (\dot{e}_{nuc}) in the burning region of each simulation. For this purpose, we define the mass-weighted average $\langle Q \rangle_w$ of some quantity Q to be

$$\langle Q \rangle_w \equiv \frac{\sum_{c_i} m(c_i) Q(c_i)}{\sum_{c_i} m(c_i)}; \quad c_i \in C_{99}(Q). \quad (1)$$

Here, $C_{99}(Q)$ is the set of grid cells with Q values in the top percentile, $Q(c_i)$ is the value of Q in cell c_i , and $m(c_i)$ is the total mass in cell c_i . Using $\langle Q \rangle_w$ instead of simply taking the maximum of the quantity across the entire simulation domain allows us to track changes over the domain as a whole rather than at a single localized point. This will therefore be a better reflection of the overall behavior of the flame rather than of a single localized fluctuation. Figure 3 shows $\langle T \rangle_w$ and $\langle \dot{e}_{\text{nuc}} \rangle_w$ as functions of time for the subset of our runs that achieve a propagating flame. This figure is referenced throughout the subsequent sections.




Figure 3. Estimates of temperature (left panel) and nuclear energy generation rate (right panel) in the burning region as functions of time. The quantities on the vertical axes are the mass-weighted averages defined in Equation 1.

3.1. Effect of Rotation Rate on Flame Structure

We run three models (F250, F500, and F1000) with the same initial model in terms of temperature but differing rotation rates. We saw in Eiden et al. (2020) that increasing the neutron star rotation rate reduces the horizontal lengthscale of the flame. An estimate of this lengthscale is given by the Rossby radius of deformation, L_R . The Rossby radius may be thought of as the scale over which the balance between the Coriolis force and horizontal pressure gradient becomes important, and is approximately given by

$$L_R \approx \frac{\sqrt{gH_0}}{\Omega}, \quad (2)$$

where g is the gravitational acceleration, H_0 is the atmospheric scale height, and Ω is the neutron star rotation rate. If we take $g \sim 1.5 \times 10^{14} \text{ cm s}^{-2}$ and $H_0 \sim 10^3 \text{ cm}$, then for the 500 Hz system we find the Rossby length to approximately be $L_R \sim 1.2 \times 10^5 \text{ cm}$. This estimate is likely to be an overestimate of the true value of the Rossby length in our simulations, as we saw that the flame is confined on a much smaller length scale. In Figure 4 and Figure 5, we use \dot{e}_{nuc} measured at 50 ms and 100 ms to discern the horizontal extent of the flame at different rotation rates. Taking the edge at greatest radius of the bright teal/green region where the most significant energy generation is occurring as the leading edge of the flame in each plot, we see that the horizontal extent of the 1000 Hz flame (F1000) appears to be reduced compared to the lower rotation 500 Hz run (F500). As the natural aspect ratio makes it hard to see the flame structure, we show the energy generation rate for the F1000 simulation with the vertical extent stretched in Figure 6. From Equation 2, we can see that increasing the rotation rate from 500 Hz to 1000 Hz should decrease L_R by a factor of two, and that the greater confinement from the Coriolis force should reduce the horizontal extent of the flame by a similar factor. However, the Rossby radius is only an approximate measure of this horizontal lengthscale, and in our simulations we see that this scaling does not work so well for all rotation rates. The simulations seem to follow the theoretical scaling described in Equation 2 more closely at higher rotation rates (1000 Hz and higher), based on the results of Eiden et al. (2020).




Figure 4. Time series of the 500 Hz run F500 showing the nuclear energy generation rate, $\dot{\epsilon}_{\text{nuc}}$.




Figure 5. Time series of the 1000 Hz run F1000 showing the nuclear energy generation rate, $\dot{\epsilon}_{\text{nuc}}$.




Figure 6. Energy generation rate for run F1000 at 100 ms with the vertical extent stretched to show detail.




Figure 7. Time series of the 500 Hz run F500 showing the mean molecular weight, \bar{A} .




Figure 8. Time series of the 1000 Hz run F1000 showing the mean molecular weight, \bar{A} .

The F500 and F1000 runs both qualitatively resemble the flame structure in Eiden et al. (2020) — a laterally propagating flame that is lifted off of the bottom of the fuel layer — but they differ in their burning structures. Figures 7 and 8 show time series of the mean molecular weight, \bar{A} , for the F500 and F1000 runs. Compared to those in Eiden et al. (2020), ashes behind the flame do not reach as high atomic weights. This is not surprising, since those previous runs artificially boosted the reaction rates. Comparing these two new runs, the burning is much more evolved for the higher rotation rate, and the ash is actually able to move ahead of the flame front (visible in the Figure 8 100 ms snapshot). We believe that this is because the increased rotation better confines the initial perturbation and subsequent expansion from the burning, increasing the temperature and density in the flame front such that the reaction rate increases, which allows the reactions to progress further. The \dot{e}_{nuc} plots in Figure 5 also support this interpretation, with the region of the flame front nearest to the crust in the F1000 run reaching higher \dot{e}_{nuc} values than for the F500 run in Figure 4. In contrast to F500 and F1000, the lowest rotation run — F250 — failed to ignite. The lack of ignition for F250 also aligns with the reasoning given above, with the lower rotation in this case potentially leading to insufficient confinement such that the temperature and density required for ignition is not achieved. In this scenario, another source of confinement (e.g. magnetic fields, see Cavecchi et al. (2016)) would need to take over at lower rotation rates to allow a burst to occur, at least for the initial model used here. Given that the size of our domain is $\sim 1 L_R$ for F250 (using Equation 2), it is also possible that we simply cannot confine the flame sufficiently with our current domain width. We see in Figure 12 (discussed further in Section 3.2) that the F500 flame took longer to achieve steady propagation than the F1000 flame. It may therefore also be that we did not run our simulation for long enough to see the F250 flame achieve the conditions required for ignition and steady propagation.

Burning in the F500 and F1000 runs is concentrated in a dense region with circular motion. In Figure 9, which compares the horizontal r -velocity u , density ρ and the nuclear energy generation rate \dot{e}_{nuc} for the F250, F500, and F1000 runs, most of the burning for each of the simulations occurs in a high density region $\rho > 3 \times 10^5 \text{ g cm}^{-3}$. The fluid in this dense, high energy generation region undergoes vortical motion, shown in the Figure 10 phase plots comparing u , the vertical z -velocity v and \dot{e}_{nuc} . This most likely corresponds to the leading edge of the flame where fresh fuel is being entrained. This feature is not developed in the 250 Hz flame in Figure 10 (left panel); it could potentially develop at later times (past the point at which we terminated our simulation), or the burning could just fizzle out and the flame fail to ignite entirely.

The mean molecular weight \bar{A} within each of our simulations seems to grow along defined tracks confined to certain temperatures T , as shown in the Figure 11 phase plots. We believe that the tracks in the plot correspond to different burning trajectories in phase space resulting from different thermodynamic condi-




Figure 9. Phase plot of the horizontal r -velocity u , density ρ and the nuclear energy generation rate \dot{e}_{nuc} for the 250, 500 and 1000 Hz runs (F250, F500 and F1000) at $t = 100$ ms. The slightly lower \dot{e}_{nuc} values along the $u = 0$ axis are most likely a numerical artifact related to the density gradient setup and finite resolution of these simulations.




Figure 10. Phase plot of the horizontal r -velocity u , vertical z -velocity v and the nuclear energy generation rate \dot{e}_{nuc} for the 250, 500 and 1000 Hz runs (F250, F500 and F1000) at $t = 100$ ms.

tions. Comparing Figure 11 to Eiden et al. (2020), these tracks are much more neat and clearly defined. The “messiness” of the tracks may be dependent on how mixed the flame interior is. Since these new simulations are un-boosted, they may be inherently less mixed than those in Eiden et al. (2020). F1000 aligns with this interpretation: its \bar{A} tracks are somewhat disrupted compared to the slower rotation runs, possibly due to the more vigorous mixing of the vortex at the flame front. Comparing the different runs, we also see that as the rotation rate increases, so does the peak temperature. This makes sense if higher rotation leads to a more concentrated, intense vortex near the flame front. It also agrees with our earlier interpretation of the enhanced burning seen in Figure 8 for F1000.

3.2. Effect of Rotation Rate on Flame Propagation

For the purpose of measuring the flame propagation speed and acceleration, we track the position of each of our flames as a function of time. We define the position in terms of a specific value of the energy generation rate, \dot{e}_{nuc} , as we did in Eiden et al. (2020). To recapitulate: we first reduce the 2D \dot{e}_{nuc} data for each simulation run to a set of 1D radial profiles by averaging over the vertical coordinate. After averaging, we take our reference \dot{e}_{nuc} value to be some fraction of the global \dot{e}_{nuc} maximum across all of these profiles.




Figure 11. Phase plot of the mean molecular weight \bar{A} , temperature T and the nuclear energy generation rate \dot{e}_{nuc} for the 250, 500 and 1000 Hz runs (F250, F500 and F1000) at $t = 100$ ms.

Table 2. Flame speed and acceleration values measured for each simulation. v_0 is the flame velocity in the $+r$ direction at $t = 0$ ms and a is the acceleration of the flame. The initial flame velocities and accelerations are derived from a quadratic least-squares fit to each of the datasets for times $t \gtrsim 20$ ms. Using these fit parameters we calculate the velocity at $t = 50$ ms, v_{50} , at which point the flames have reached steady propagation.

run	v_0 (km s $^{-1}$)	a (km s $^{-2}$)	v_{50} (km s $^{-1}$)
F1000	3.414 ± 0.008	-1.03 ± 0.13	3.36 ± 0.01
F500	3.077 ± 0.013	5.41 ± 0.19	3.35 ± 0.02
F500_2E8	3.760 ± 0.015	1.74 ± 0.25	3.85 ± 0.02
F500_3E8	-5.293 ± 0.100	357.41 ± 2.15	12.58 ± 0.15

Since the flames in our simulations propagate in the positive horizontal direction, we then search the region of each profile at greater radius than the local \dot{e}_{nuc} maximum for the point where the \dot{e}_{nuc} first drops below this reference value. This point gives us the location of our flame front.

In Eiden et al. (2020), we used 0.1% of the global \dot{e}_{nuc} maximum for our reference value. For the high temperature unboosted flames, however, we found that the \dot{e}_{nuc} profiles failed to reach that small a value across the domain at most times, which prevented us from obtaining reliable position measurements. We therefore use 1% of the global \dot{e}_{nuc} maximum in this paper rather than 0.1%. This is still sufficiently small that our measurements are not overly sensitive to turbulence and other local fluid motions (the issue with simply tracking the local maximum), but allows us to avoid the pitfall encountered by the 0.1% metric.




Figure 12. Flame front position vs. time for the standard (10^8 K) 500 Hz and 1000 Hz runs (F500 and F1000). The dashed lines show quadratic least-squares fits to the data for $t \gtrsim 20$ ms.

Figure 12 gives the radial position of the flame front as a function of time for the F500 and F1000 runs (blue and orange, respectively) to show the dependence on rotation rate. In Eiden et al. (2020), we applied a linear least-squares fit to the flame front position as a function of time to estimate the propagation velocity. As some of the flames in this set of simulation runs exhibit significant acceleration, for this study we instead fit the data with a quadratic curve of the form

$$r(t) = \frac{1}{2}at^2 + v_0t + r_0, \quad (3)$$

where the parameter a is the acceleration of the flame, v_0 is the velocity at $t = 0$ ms, and r_0 is the flame front position at $t = 0$ ms. We do not include the data points with $t \lesssim 20$ ms when performing the fit, since these correspond to the initial transient period before the flame has begun to propagate steadily. The values of a and v_0 for the full suite of simulation runs are provided in Table 2. Note that v_0 is only a parameter that may be used to calculate the flame speed at an arbitrary time. It is not an estimate of the true initial velocity of the flame, since the flame has not achieved ignition yet at $t = 0$ ms. We use the fit parameters to calculate the flame speeds at $t = 50$ ms (when the flame has reached steady propagation), given in the fourth column of Table 2.

For the simple case of a laminar flame front driven by conduction, the propagation rate, s_l , of the flame is approximately

$$s_l \approx \left(\frac{\mathcal{D}\dot{\epsilon}_{\text{muc}}}{E} \right)^{\frac{1}{2}}. \quad (4)$$

Here, $\mathcal{D} = k_{th}/(\rho c_v)$ is the thermal diffusivity with thermal conductivity k_{th} and specific heat capacity c_v , \dot{e}_{nuc} is the specific nuclear energy generation rate, and E is the specific total energy (Timmes & Woosley 1992). Cavecchi et al. (2013) explored the effects of rotation on the flame’s propagation, and noted that the flame speed in their simulations scaled with the ratio L_R/H_0 of Rossby length to atmospheric scale height. Combining this scaling with Equation 4 yields a revised estimate for the flame speed, s :

$$s \approx \frac{L_R}{H_0} s_l \approx \frac{L_R}{H_0} \left(\frac{\mathcal{D} \dot{e}_{nuc}}{E} \right)^{\frac{1}{2}}. \quad (5)$$

Observing that L_R is inversely proportional to the angular rotation frequency Ω and neglecting variations in H_0 and E , we can then derive the following approximate scaling relation:

$$s \propto \frac{(\mathcal{D} \dot{e}_{nuc})^{\frac{1}{2}}}{\Omega}. \quad (6)$$

We were able to test the validity of this scaling relation in Eiden et al. (2020), where we studied the effects of varying the thermal conductivity/diffusivity, energy generation rate, and rotation rate. These results supported the findings of Cavecchi et al. (2013). Despite this formulation of the flame speed neglecting some relevant physics (e.g. turbulence), the relation proved consistent with our simulation results.

This study explores a more physical parameter space than Eiden et al. (2020), and as seen in Figure 12, there is no clear inverse scaling of the flame speed with rotation rate in the new set of runs. We observed earlier, however, that nuclear reactions progress more quickly at higher rotation rate. This results in a higher \dot{e}_{nuc} — up to three to four times higher near the burning front after the flame ignites (see Figure 3) — which may counteract the reduction in flame speed from the increased Coriolis confinement. Comparing accelerations, we also observe that F500 accelerates faster than F1000, which appears to experience a small deceleration at early times. This disparity may be a direct result of the difference in Coriolis force.

3.3. Effect of Temperature on Flame Structure

To explore the effect of different initial temperature configurations, we run four simulations fixed at a rotation rate of 500 Hz with temperatures as shown in Table 1. For all the 500 Hz simulations (with the exception of the coolest run, F500), we set $T_{star} = T_{hi}$, scaling δT accordingly to maintain a consistent value of $T_{hi} + \delta T$. If we let $T_{star} < T_{hi}$, the cooler neutron star surface might act as a heat sink and siphon away energy that would otherwise go into heating the burning layer. By choosing $T_{star} = T_{hi}$, we can effectively explore simulations with greater surface heating. There are several physically distinct mechanisms that could produce an increased temperature at the crust: crustal heating, some other form of shallow heating or accretion-induced heating. In these simulations, we do not model the mechanism producing the heating effect, just the effect itself, so we do not distinguish between which of these mechanisms cause the heating.

Figure 13 shows \bar{A} for three 500 Hz simulations with different initial temperature structures ($T_{star} \leq 3 \times 10^8$ K) at $t = 70$ ms. We do not plot F500.4E8 here because it fails to form a clear burning front. F500.3E8 (Figure 13, bottom panel) — the hottest run to form a clear burning front — has a faster propagating flame (this will be discussed further in Section 3.4). It also reaches slightly higher \bar{A} values than the two cooler runs. The Figure 19 \bar{A} - T phase plots of F500 (left) and F500.3E8 (middle) also reflect these \bar{A} features, with F500.3E8 reaching slightly higher \bar{A} values. F500.3E8 also reaches higher \dot{e}_{nuc} values at the low end of the temperature range ($< 0.5 \times 10^9$ K). There appear to be more causally connected regions across a range of \bar{A} at low temperatures for F500.3E8 than for F500, suggesting that the higher \dot{e}_{nuc} for F500.3E8




Figure 13. Comparison of \bar{A} for 3 different 500 Hz models with neutron star temperatures T_{star} of 10^8 K, 2×10^8 K and 3×10^8 K (F500, F500_2e8, and F500_3e8, respectively), and resulting envelope structures. Each flame is shown at 70 ms.

generates burning in certain burning trajectories that are not present in the cooler F500 run. Note that, although F500_3E8 is the hottest run with a modified initial temperature configuration to form a distinct flame front, the highest rotation F1000 flame actually reaches higher temperatures (see Figure 3, left panel) as well as higher \bar{A} (see Figure 8).

In contrast to the models with $T_{\text{star}} \leq 3 \times 10^8$ K, F500_4E8 is so hot that the organized flame structure is lost. This model burns so strongly that it is only run out to 40 ms. After an initial period when the burning moves across the domain, residual burning continues and eventually ignites the entire fuel layer at late times, as shown in Figures 14 and 15 for three snapshots taken in the last ten seconds of the simulation. In Figure 14, \dot{e}_{nuc} reaches values of $10^{18} - 10^{20}$ erg g $^{-1}$ s $^{-1}$ across the domain and at heights up to $\sim 0.5 \times 10^4$ cm. There is still significant burning occurring even higher, with \dot{e}_{nuc} reaching $\sim 10^{15}$ erg g $^{-1}$ s $^{-1}$ at heights up to $\sim 0.8 \times 10^4$ cm. For comparison, the next hottest run (F500_3E8) only reaches maximum \dot{e}_{nuc} values on the order of 10^{18} erg g $^{-1}$ s $^{-1}$ (see Figure 3, right panel), even at the latest timesteps (~ 70 ms, when the flame is most developed). Significant \dot{e}_{nuc} values for all runs other than F500_4E8 are confined to the flame front, rather than spanning the entire domain.

Rather than the fuel igniting in a localized hot spot and forming a flame which propagates across the domain, for this system the heat produced by the high temperature crust alone is sufficient to ignite material across the entire domain. Uniform burning across a large fraction of or even the entire surface of the neutron star has been observed in neutron stars with accretion rates that exceed the Eddington limit (Fujimoto et al. 1981; Bildsten 1998; Keek et al. 2009). However, it is important to note here that the conditions in our model vary from the conditions that would occur in a neutron star ocean undergoing stable burning in a number of ways. Firstly, the model begins without any burning taking place, and the initial model assumes that a significant amount of unburnt fuel has been able to build up. For a stably burning system, freshly accreted material would be continuously burnt, so there would be much less unburnt fuel available. Secondly, we are only modelling a few tens of milliseconds, rather than an ongoing process. We therefore do not demonstrate that this system would continue to burn stably if the high crustal temperature were maintained and sufficient fuel continued to be available, e.g. through the continued accretion of fresh material. However, this run does suggest that a crust with a sufficiently high temperature can provide enough energy to heat the fuel above it to the point of ignition, and it may be that if these conditions were maintained that it would be able to sustain stable burning.




Figure 14. Time series of the 500 Hz run F500_4E8, with $T_{\text{star}} = 4 \times 10^8$ K, showing \dot{e}_{nuc} . This model burns so strongly that it is only run out to 40 ms; the snapshots shown here are at $T = 30$ ms, 35 ms, and 40 ms.

Figure 15 shows \bar{A} for F500_4E8. Again, burning extends across the domain and high into the atmosphere and fuel layer, lacking the characteristic flame structure shown in \bar{A} plots for the lower temperature 500 Hz runs (see Figure 13). A distinct mass of material appears to have broken out of the atmosphere near the axis of symmetry. The atmosphere edge is located at $\sim 1.2 \times 10^4$ cm. A similar effect is also visible (to a lesser extent) in the F500_3E8 plot in Figure 13, with a faint haze of material rising above the flame near the axis of symmetry. However, this is likely to be a numerical artefact rather than a true physical effect, and a result of the boundary conditions imposed at the axis of symmetry for these simulations. F500_4E8 clearly reaches much higher \bar{A} values across the domain (especially at the latest snapshot, $t = 40$ ms) compared to all the other runs described in this paper (see Figures 7, 8 and 13).

Though F500_4E8 does not form a distinct burning front, it does achieve greatly enhanced burning, as shown in the Figure 19 \bar{A} - T phase plot (right panel). Not only is F500_4E8 significantly hotter overall compared to all other runs, but there is a large causally connected region across a wide range of \bar{A} . This indicates that the hotter temperature has facilitated significant burning in burning trajectories that were not favored at the lower temperatures. The burning trajectories are also very disrupted for F500_4E8 compared to the cooler runs, suggesting that the hotter temperature leads to more vigorous mixing. Indeed, this appears to be the case looking at Figure 15 and the right panel of Figure 18 (discussed further in section 3.4). The disrupted burning trajectories resemble those found in Figure 11 for the highest rotation F1000 run, though they are even more dramatically disrupted for F500_4E8. F500_4E8 is also noticeably hotter than F1000, even though it is run for significantly less time ($t = 40$ ms vs $t = 100$ ms). Although the F500_4E8 run is clearly a special case in that it develops steady burning across the domain rather than a propagating flame, enhanced burning in this hottest run aligns with results from the other simulations with differing initial temperature structures. As T_{star} is increased, the overall behavior and propagation of the flame is significantly altered, implying that burning is very temperature-sensitive. We explore flame propagation further in Section 3.4.




Figure 15. Time series of the 500 Hz run F500_4E8, with $T_{\text{star}} = 4 \times 10^8$ K, showing \bar{A} .

3.4. Effect of Temperature on Flame Propagation

The method for measuring flame propagation described in Section 3.2 is applied in Figure 16 to the three 500 Hz runs with $T_{\text{star}} \leq 3 \times 10^8$ K. Due to the lack of a clear burning front in F500_4E8, we do not analyze its propagation velocity and acceleration. As the initial T_{star} is increased beyond $\sim 2 \times 10^8$ K, the flame becomes greatly accelerated. The initial flame velocities and accelerations derived from a quadratic least-squares fit to each of the datasets, as well as the flame velocities at $t = 50$ ms calculated using the fit parameters, are provided in Table 2. Comparing the flame propagation at different initial temperatures, the most robust feature is the acceleration of the F500_3E8 run at $t \sim 40$ ms. The reason for the acceleration of the flame is not entirely clear. Whereas for the cooler runs, a state of steady flame propagation is achieved, for the F500_3E8 run the flame speed continues to increase, suggesting that some instability driving the flame's propagation persists to later times. It could be that energy released from burning begins to dominate the flame's propagation as it evolves, increasing the flame speed over time. Another possibility is that the increased temperatures lead to enhanced turbulent mixing effects that pull in more fresh fuel for the flame to burn. Yet another possibility is that the higher initial T_{star} leads to a greater average temperature in the fuel layer over time, making it easier for the flame to burn the fuel and propagate.

In both the u - ρ phase plot in Figure 17 and the u - v phase plot in Figure 18, the horizontal u -velocity for the F500_4e8 (right panel) burning region is slightly larger than that of the cooler runs. Additionally, the cooler runs' u -velocities are two to three orders of magnitude greater than the flame speeds listed in Table 2. These phase plot velocities therefore most likely correspond to vortical motion in the turbulent burning vortices rather than the propagation of a flame itself. The vertical z -velocity v in Figure 18 further suggests that F500_4e8 undergoes increased vortical motion, as we see that the velocity magnitude in the burning region is significantly larger than in the lower temperature runs. Also of interest in these plots is that the coolest run shows high velocity material comprising both high and low \dot{e}_{nuc} with horizontal velocity asymmetry, both of which suggest the burning vortex in the coolest case is less well defined than at higher temperatures. The coolest run thus does not seem to have developed the characteristic vortex structure at




Figure 16. Flame front position vs. time for the three 500 Hz runs with T_{star} of 10^8 K, 2×10^8 K and 3×10^8 K (F500, F500_2e8, and F500_3e8, respectively). The dashed lines show quadratic least-squares fits to the data for $t \gtrsim 20$ ms. Note that due to its rapid acceleration, F500_3E8 (green) is only run out to ~ 70 ms to avoid surpassing the domain boundary.

the burning front (i.e. where \dot{e}_{nuc} is greatest) that can be clearly seen for the hotter runs at this time (40 ms). As can be seen in Figure 10 (which was plotted at 100 ms), this does develop more at later times. Similar to what we saw when comparing the runs with different rotation rates, it would therefore appear that it takes longer for the flame to develop when T_{star} is cooler.

4. DISCUSSION AND CONCLUSIONS

We ran a number of simulations of laterally propagating flames in XRBs in order to explore the effects of rotation and thermal structure. We found that increasing the rotation rate increased the energy generation rate within the flame and enhanced nuclear burning. Apart from the lowest rotation run (which failed to ignite), flame propagation was not noticeably impacted by rotation rate; by the time the different flame fronts reached steady propagation, they shared comparable velocities. These results are likely due to the rotation-dependent strength of the Coriolis force and its confinement of the flame balancing the enhanced nuclear burning. This lack of dependence of the flame speed on the rotation rate could be tested using observations, as it implies that the burst rise time should have little dependence on the rotation rate.

We explored several models with different crustal temperatures to determine what effects mechanisms such as higher accretion rates, crustal heating and shallow heating may have on flame propagation. We found that increasing the temperature of the crust significantly enhanced the flame propagation. This we believe to be because a cooler crust allows heat to more efficiently be transferred away from the flame itself, therefore reducing the flame's temperature, slowing burning and consequently reducing its propagation




Figure 17. Phase plot of the horizontal r -velocity u , density ρ and the nuclear energy generation rate \dot{e}_{nuc} for the 500 Hz runs with T_{star} of 10^8 K, 3×10^8 K and 4×10^8 K (F500, F500_3e8, and F500_4e8) at $t = 40$ ms (the latest time that F500_4E8 is run out to). The run with $T_{\text{star}} = 2 \times 10^8$ K (F500_2E8) closely resembles the cooler F500 run, and is not shown here.




Figure 18. Phase plot of the horizontal r -velocity u , vertical z -velocity v and the nuclear energy generation rate \dot{e}_{nuc} for the 500 Hz runs with T_{star} of 10^8 K, 3×10^8 K and 4×10^8 K (F500, F500_3e8, and F500_4e8) at $t = 40$ ms. The run with $T_{\text{star}} = 2 \times 10^8$ K (F500_2E8) closely resembles the cooler F500 run, and is not shown here.

speed. At higher crustal temperatures, we saw that the inability for heat to be efficiently transported away from the flame front increased the flame temperature, driving unstable, accelerating flame propagation. We saw that if the crust temperature was too high, then instead of a flame the entire atmosphere would burn steadily. This is reminiscent of what is seen for neutron stars with accretion rates exceeding the Eddington limit. This dependence of the flame speed on the crustal temperature could also be tested using observations, by looking to see how the burst rise time depends on the accretion rate (which may affect the crustal heating).

In future work, we would like to improve and expand our simulations in order to better understand the processes at play and to include more physics. This includes adding tracer particles to the simulations so we can monitor the fluid motion and perform more detailed nucleosynthesis; extending our simulations to 3D, which would hopefully alleviate some of the boundary effects we have observed in these simulations but will require significant computational resources; and exploring the resolution of our simulations more so that we can ensure that we have resolved all of the necessary physical processes. We would also like to model H/He flames, as these are the sites of rp-process nucleosynthesis (Schatz et al. 2001). Initially we




Figure 19. Phase plot of the mean molecular weight \bar{A} , temperature T and the nuclear energy generation rate $\dot{\epsilon}_{\text{nuc}}$ for the 500 Hz runs with T_{star} of 10^8 K, 3×10^8 K and 4×10^8 K (F500, F500_3e8, and F500_4e8) at $t = 40$ ms. The run with $T_{\text{star}} = 2 \times 10^8$ K (F500_2E8) closely resembles the cooler F500 run, and is not shown here.

will use the same reaction sequence explored in our previous convection calculations (Malone et al. 2014). Finally, we recently added an MHD solver to **Castro** (Barrios Sazo 2020); this will allow us in the future to explore the effects of magnetic fields on flame propagation in XRBs.

ACKNOWLEDGMENTS

Castro is open-source and freely available at <http://github.com/AMReX-Astro/Castro>. The problem setup used here is available in the git repo as `flame.wave`. We thank the anonymous referee for their helpful suggestions. The work at Stony Brook was supported by DOE/Office of Nuclear Physics grant DE-FG02-87ER40317. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research and Office of Nuclear Physics, Scientific Discovery through Advanced Computing (SciDAC) program under Award Number DE-SC0017955. This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. This material is also based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0021110. This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Science Undergraduate Laboratory Internship (SULI) program. MZ acknowledges support from the Simons Foundation. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U. S. Department of Energy under Contract No. DE-AC02-05CH11231. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725, awarded through the DOE INCITE program. We thank NVIDIA Corporation for the donation of a Titan X and Titan V GPU through their academic grant program. This research has made use of NASA’s Astrophysics Data System Bibliographic Services.

Facilities: NERSC, OLCF

Software: AMReX (Zhang et al. 2019), Castro (Almgren et al. 2010; Almgren et al. 2020), GCC (<https://gcc.gnu.org/>), Jupyter (Kluyver et al. 2016), linux (<https://www.kernel.org/>), matplotlib (Hunter 2007, <http://matplotlib.org/>), NumPy (Oliphant 2007; van der Walt et al. 2011), python (<https://www.python.org/>), valgrind (Nethercote & Seward 2007), VODE (Brown et al. 1989), yt (Turk et al. 2011)

REFERENCES

- Almgren, A., Sazo, M. B., Bell, J., et al. 2020, Journal of Open Source Software, 5, 2513, doi: [10.21105/joss.02513](https://doi.org/10.21105/joss.02513)
- Almgren, A. S., Beckner, V. E., Bell, J. B., et al. 2010, ApJ, 715, 1221, doi: [10.1088/0004-637X/715/2/1221](https://doi.org/10.1088/0004-637X/715/2/1221)
- Altamirano, D., Watts, A., Kalamkar, M., et al. 2010, ATel, 2932, 1
- Barrios Sazo, M. G. 2020, PhD thesis, State University of New York at Stony Brook
- Bernard-Champmartin, A., Braeunig, J.-P., & Ghidaglia, J.-M. 2012, Computers and Fluids, 7, doi: [10.1016/j.compfluid.2012.09.014](https://doi.org/10.1016/j.compfluid.2012.09.014)
- Bhattacharyya, S., & Strohmayer, T. E. 2007, 666, L85, doi: [10.1086/521790](https://doi.org/10.1086/521790)
- Bildsten, L. 1998, in ASIC, Vol. 515, 419
- Bilous, A. V., & Watts, A. L. 2019, 245, 19, doi: [10.3847/1538-4365/ab2fe1](https://doi.org/10.3847/1538-4365/ab2fe1)
- Brown, P. N., Byrne, G. D., & Hindmarsh, A. C. 1989, SIAM J. Sci. Stat. Comput., 10, 1038
- Cavecchi, Y., Levin, Y., Watts, A. L., & Braithwaite, J. 2016, MNRAS, 459, 1259, doi: [10.1093/mnras/stw728](https://doi.org/10.1093/mnras/stw728)
- Cavecchi, Y., Watts, A. L., Braithwaite, J., & Levin, Y. 2013, MNRAS, 434, 3526, doi: [10.1093/mnras/stt1273](https://doi.org/10.1093/mnras/stt1273)
- Cavecchi, Y., Watts, A. L., Levin, Y., & Braithwaite, J. 2015, MNRAS, 448, 445, doi: [10.1093/mnras/stu2764](https://doi.org/10.1093/mnras/stu2764)

- Chakraborty, M., & Bhattacharyya, S. 2014, *ApJ*, 792, 4, doi: [10.1088/0004-637X/792/1/4](https://doi.org/10.1088/0004-637X/792/1/4)
- Cumming, A., Macbeth, J., Zand, J. J. M. i. T., & Page, D. 2006, *The Astrophysical Journal*, 646, 429, doi: [10.1086/504698](https://doi.org/10.1086/504698)
- Deibel, A., Cumming, A., Brown, E. F., & Page, D. 2015, *ApJL*, 809, L31, doi: [10.1088/2041-8205/809/2/L31](https://doi.org/10.1088/2041-8205/809/2/L31)
- Eiden, K., Zingale, M., Harpole, A., et al. 2020, *ApJ*, 894, 6, doi: [10.3847/1538-4357/ab80bc](https://doi.org/10.3847/1538-4357/ab80bc)
- Fujimoto, M. Y., Hanawa, T., & Miyaji, S. 1981, *Astrophysical Journal*, 247, 267, doi: [10.1086/159034](https://doi.org/10.1086/159034)
- Galloway, D. K., in 't Zand, J. J. M., Chenevez, J., et al. 2020, arXiv e-prints, arXiv:2003.00685. <https://arxiv.org/abs/2003.00685>
- Hunter, J. D. 2007, *Computing in Science and Engg.*, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- Johnston, Z., Heger, A., & Galloway, D. K. 2019, arXiv e-prints, arXiv:1909.07977. <https://arxiv.org/abs/1909.07977>
- Katz, M. P., Almgren, A., Sazo, M. B., et al. 2020, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '20* (IEEE Press)
- Keek, L., & Heger, A. 2017, *The Astrophysical Journal*, 842, 113, doi: [10.3847/1538-4357/aa7748](https://doi.org/10.3847/1538-4357/aa7748)
- Keek, L., Langer, N., et al. 2009, *Astronomy & Astrophysics*, 502, 871
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, 87–90, doi: [10.3233/978-1-61499-649-1-87](https://doi.org/10.3233/978-1-61499-649-1-87)
- Malone, C. M., Zingale, M., Nonaka, A., Almgren, A. S., & Bell, J. B. 2014, *ApJ*, 788, 115, doi: [10.1088/0004-637X/788/2/115](https://doi.org/10.1088/0004-637X/788/2/115)
- Nethercote, N., & Seward, J. 2007, in *Proceedings of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '07* (New York, NY, USA: ACM), 89–100, doi: [10.1145/1250734.1250746](https://doi.org/10.1145/1250734.1250746)
- Oliphant, T. E. 2007, *Computing in Science and Engg.*, 9, 10, doi: [10.1109/MCSE.2007.58](https://doi.org/10.1109/MCSE.2007.58)
- Schatz, H., Aprahamian, A., Barnard, V., et al. 2001, *Physical Review Letters*, 86, 3471, doi: [10.1103/PhysRevLett.86.3471](https://doi.org/10.1103/PhysRevLett.86.3471)
- Spitkovsky, A., Levin, Y., & Ushomirsky, G. 2002, *ApJ*, 566, 1018, doi: [10.1086/338040](https://doi.org/10.1086/338040)
- Timmes, F. X. 2000, *ApJ*, 528, 913, doi: [10.1086/308203](https://doi.org/10.1086/308203)
- Timmes, F. X., & Woosley, S. E. 1992, *ApJ*, 396, 649, doi: [10.1086/171746](https://doi.org/10.1086/171746)
- Turk, M. J., Smith, B. D., Oishi, J. S., et al. 2011, *ApJS*, 192, 9, doi: [10.1088/0067-0049/192/1/9](https://doi.org/10.1088/0067-0049/192/1/9)
- Turlione, A., Aguilera, D. N., & Pons, J. A. 2015, *Astronomy and Astrophysics*, 577, doi: [10.1051/0004-6361/201322690](https://doi.org/10.1051/0004-6361/201322690)
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *Computing in Science & Engineering*, 13, 22, doi: [10.1109/MCSE.2011.37](https://doi.org/10.1109/MCSE.2011.37)
- Zhang, W., Almgren, A., Beckner, V., et al. 2019, *Journal of Open Source Software*, 4, 1370, doi: [10.21105/joss.01370](https://doi.org/10.21105/joss.01370)