Phase-Aligned Fuel Injection in a Reduced Burn Propagation Model

Martin Chatwood*
ORCiD: https://orcid.org/0009-0009-0982-2290
Chatwood Labs Ltd, United Kingdom
(Dated: April 07, 2026)

Abstract

Fuelling studies in magnetically confined plasmas typically focus on fuelling rate, deposition location, and penetration depth. The temporal phase of fuelling relative to evolving burn structures has received far less attention. In systems where burn propagates as a moving structure, timing may also influence the outcome. This work examines that possibility using a reduced one-dimensional toroidal model in which crest propagation is represented through prescribed advection together with simplified source, transport, thermalisation, and exhaust closures. The model is intended as a conceptual mechanism study rather than a predictive plasma calculation.

The model supports a travelling burn crest sustained by fusion heating and localised fuel input. Injecting fuel at different angular offsets relative to the tracked crest reveals strong phase sensitivity in the lag scan. Injection close to the crest produces the strongest local burn response, but this regime is short-lived. By contrast, delayed injection into the wake sector sustains propagation for much longer, whereas asynchronous injection collapses on nearly the same timescale as the no injection baseline despite receiving the same scheduled fuel input per orbit as the phase-aligned cases.

Sensitivity tests confirm that the behaviour remains robust across a range of parameters, with grid refinement demonstrating convergence of the crest structure and thermalisation sweeps preserving the same qualitative lag dependence. To probe the model more broadly, two Monte Carlo ensembles were generated: one varying thermalisation and pumping parameters, and another varying crest geometry and transport parameters around baseline. Across these runs, the phase-dependent propagation window persists. The response shows a clear transition between a strong-burn, short-lived near-crest regime and a longer-lived delayed-fuelling regime, with the longest runtimes occurring in the sector closer to the anti-crest side of the domain. Across the tested parameter sets, this longer-lived phase window persists with only modest shifts in location, while termination-criterion variation mainly affects the detected collapse boundary rather than the existence of the regime.

The results indicate that structured fuelling can influence propagation dynamics in this reduced system even when total fuel input is unchanged. The model is intentionally simplified and excludes magnetic geometry and stability physics. It is presented as an advection-diffusion-reaction dynamics study of phase-coupled forcing, not a plasma prediction, and motivates further investigation of phase-structured fuelling in higher-fidelity simulations.

1. Introduction

Fuel injection is normally discussed in terms of magnitude. Increasing the fuel supply, improving confinement, or raising heating power are the usual levers considered when analysing burn behaviour in fusion plasmas. These factors determine whether a plasma reaches ignition and how a burning state can be sustained. Much of the modern understanding of burning plasma behaviour is framed through global energy

* research@chatwoodlabs.com

balance, transport scaling, and confinement performance, as summarized in the ITER physics basis and subsequent updates that describe the expected operational regimes of burning plasmas [1] [2].

In most analyses the burn state is treated as a largely uniform condition of the plasma column. Even when spatial structure is considered, the dominant emphasis remains on energy balance and transport rather than the timing of fuel delivery relative to evolving local structures. This perspective is natural in reactor design where integrated plasma performance, stability, and confinement constraints dominate the discussion.

If burn evolves as a travelling structure rather than a uniform state, the timing of fuel delivery may influence its evolution. A propagating crest can interact differently with fuel introduced ahead of it, within it, or in its wake. In such systems the phase relationship between a propagating burn front and localised fuel input becomes a potentially important control parameter.

Burn propagation as a spatially structured front rather than a uniform state has precedent in fusion literature. Christopherson, A. R. [3] developed analytic theory of burn propagation in inertially confined plasmas, demonstrating that thermonuclear burn extends from a localised hotspot into surrounding cold fuel once alpha heating drives ignition. Related analyses of inertial confinement fusion, including Lindl [4], describe the role of localised ignition in initiating burn that spreads into colder fuel. A more complete theoretical framework is provided by Atzeni and Meyer-ter-Vehn [5], which describes the development of thermonuclear burn waves sustained by alpha particle self-heating. Although these treatments concern inertial-fusion-related regimes, they show that spatially advancing thermonuclear burn regions are physically admissible in localised-fusion parameter regimes.

Propagation of thermonuclear burn fronts has been examined in more recent inertial confinement contexts as well. Burn waves can propagate through magnetised or compressed fuel once local heating and fusion reactions become self-sustaining. Recent work has also examined burn propagation in magnetised inertial-fusion settings, showing that localised heating and fuel distribution can influence burn evolution [6].

In magnetic confinement, a coherent toroidally propagating thermonuclear burn structure has not been experimentally established, although burn propagation has been discussed conceptually in the burning-plasma literature, particularly in high-Q tokamak contexts. Here, that structure is adopted as a modelling assumption in order to ask whether the phase relationship between localised fuel injection and a moving front influences its persistence. This is a question about propagation dynamics rather than ignition physics and is examined in a reduced framework.

Direct evidence for propagating thermonuclear burn fronts in magnetically confined plasmas is lacking. However, several classes of propagating plasma phenomena are well established. These include energetic-particle-driven Alfvénic activity [7], heat-pulse propagation following localised perturbations [8], and turbulence spreading in drift-wave systems [9]. These phenomena are not burn waves, but they do show that magnetised plasmas can sustain coherent, evolving structures governed by coupled transport and energy dynamics. Phase-dependent interactions between localised fuelling and a propagating structure are therefore not inherently implausible, even if direct toroidal burn-front analogues have not been demonstrated.

More generally, phase relationships play a central role in many nonlinear dynamical systems. Synchronisation and phase locking phenomena appear in coupled oscillator populations, fluid systems, and a wide range of physical and biological networks. Classical work on synchronisation shows how systems with many interacting components can transition between incoherent and phase-locked states depending on coupling strength and phase alignment [10].

The present work examines whether a related phase sensitivity can appear in a reduced model of burn propagation. A reduced toroidal system is constructed in which temperature, fuel density,

and transport evolve along a one-dimensional periodic domain. Such reduced nonlinear transport models are commonly used to study propagating structures in nonequilibrium systems [11]. Fusion heating, thermalisation of injected fuel, and exhaust losses are represented through simplified source and transport terms. Within this framework a travelling temperature crest forms and propagates along the toroidal coordinate. The instantaneous crest position is defined as the toroidal location of the maximum temperature at each timestep.

Fuel can be injected at a controlled angular offset relative to this crest. The offset is described as a wake-lag measured along the toroidal coordinate. This configuration allows the timing of fuel delivery relative to the propagating structure to be varied systematically while holding the total fuel input constant.

This study uses a reduced conceptual model that omits magnetic geometry, MHD stability, fast-particle dynamics, and other reactor-relevant effects. It asks only whether, once burn propagation is treated as a travelling structure, burn-front persistence becomes sensitive to the phase of localised fuel injection.

To investigate this question a wake-lag scan is performed in which the injection offset is varied across the toroidal domain. The results are compared with two control cases: a no injection baseline and asynchronous injection in which fuel is delivered without phase alignment. Additional tests examine grid convergence, parameter sensitivity, and robustness under moderate perturbations of transport and source parameters.

The simulations indicate that the reduced system exhibits a clear phase-sensitive response. Injection close to the crest produces the strongest local burn and the highest crest temperatures. In the baseline configuration, delayed injection induces a sharp transition from baseline-like collapse to extended propagation over a narrow interval near |lag| ≈ 64° (about 63.5°-64.5° on the positive branch at the present scan resolution), with the maximum lifetime occurring near the anti-crest regime around ±180°. In contrast, asynchronous delivery of the same total fuel input yields collapse on essentially the same timescale as the no injection baseline. Although the model is intentionally simplified, the behaviour suggests that the timing of fuel delivery relative to a propagating burn structure can influence its persistence, making phase-structured fuelling worth testing in higher-fidelity plasma simulations.

2. Reduced Propagation Model

The system studied here is a reduced conceptual representation of phase-coupled fuelling relative to a travelling burn structure along a toroidal coordinate. The plasma is represented as a one-dimensional periodic ring. In this framework, crest propagation is not derived from a self-consistent reactor-scale burn-wave calculation; it is represented through prescribed nominal advection together with simplified source, transport, thermalisation, and exhaust closures. The purpose of the model is therefore to isolate how injection phase interacts with this travelling structure, not to predict reactor plasma behaviour.

The reduced system can be written schematically in reaction-transport form as commonly used in reduced models of nonlinear transport and pattern formation [11]:

$$\frac{\partial U}{\partial t} = -v_c \frac{\partial U}{\partial \theta} + D_U \frac{\partial^2 U}{\partial \theta^2} + Q_{fusion}(n,T) - L(T) + S_{inj},$$

$$\partial \frac{\partial n}{\partial t} = -v_n \frac{\partial n}{\partial \theta} + D_n \frac{\partial^2 n}{\partial \theta^2} - \Gamma_{exhaust} + S_{inj},$$

where U represents the local plasma energy density, T the temperature field, and n the thermalised fuel density. The coefficients D_U and D_n represent effective transport coefficients for energy and particle diffusion along the toroidal coordinate. The advection terms with velocities $v_c$and $v_n$represent directed propagation of energy and fuel associated with the moving burn crest.

The term $Q_{\text{fusion}}(n,T)$represents local fusion heating produced by deuterium-tritium reactions, while $L(T)$represents temperature-dependent loss processes. The term Γ_{exhaust}represents particle

removal due to exhaust or pumping. The source term S_{inj}represents localised fuel injection events.

Temperature is obtained diagnostically from the ratio of local energy density to thermalised fuel density,

$$T = \frac{U}{n}.$$

Here U denotes the reduced thermal energy field associated with the diagnostic reacting-temperature variable, so the temperature closure is taken over the thermalised DT population only, T = U/(n_D + n_T); helium ash is evolved and exhausted separately and is not included in this heat-capacity closure. In the numerical implementation, injected fuel first enters a fresh, non-reacting reservoir and only later thermalises into the burn-eligible fuel population at a prescribed rate, while injected enthalpy is accounted for at injection time; the explicit implementation and diagnostic definitions are given in Appendix C. The resulting interaction between fusion heating, transport, fuel injection, and particle exhaust produces a travelling temperature crest that propagates along the periodic toroidal coordinate.

The domain is discretized into N cells with periodic boundary conditions. Each cell contains the local temperature and particle densities for deuterium tritium fuel and helium ash. Transport along the ring combines directed advection associated with crest propagation and diffusive smoothing of energy and particle gradients. Fusion heating provides a localised energy source that depends on the local temperature and fuel density. Fuel injection acts as an external source while exhaust removes particles from the system on specified timescales.

Within this framework a travelling temperature crest can form, analogous to propagating fronts studied in nonlinear reaction-diffusion systems [12]. The crest moves along the toroidal coordinate as heating, transport, and fuel supply interact. Simulation termination occurs when the crest loses coherence according to criteria based on crest flatness, speed, and temperature. Because crest collapse is defined diagnostically through these criteria, sensitivity tests are later performed to verify that the observed phase-dependent behaviour does not depend on a particular stopping rule.

In the simulations the crest position is defined as the angular location of the maximum temperature along the toroidal domain, identified at each timestep using an argmax operation on the temperature field. This definition is used to measure crest position, propagation speed, and the relative phase between injection events and the propagating crest.

2.1. Governing Quantities

Several coupled fields are evolved along the ring:

- Energy density U(θ,t) from which temperature T(θ,t) is derived
- Thermalised deuterium and tritium fuel densities n_D(θ,t) and n_T(θ,t)
- Fresh injected fuel densities n_D,fresh(θ,t) and n_T,fresh(θ,t)
- Helium ash density n_He(θ,t)

Fusion reactions convert fuel into helium while depositing energy locally in the plasma. The resulting heating contributes to the evolution of the temperature field. Injected fuel first enters a fresh, non-reacting reservoir and subsequently transfers into the thermalised fuel population at a prescribed thermalisation rate. Only thermalised fuel participates in fusion reactions.

The goal of this system is not to reproduce detailed fusion plasma physics. Instead, it isolates the interaction between localised fuelling and a propagating burn structure. Magnetic geometry, magnetohydrodynamic stability, and kinetic effects are not included.

2.2. Fuel injection and wake-lag

Fuel is injected into the system through a localised source term that deposits fresh deuterium-tritium fuel into the non-reacting reservoir. The injection location is defined relative to the instantaneous position of the temperature crest. This offset is expressed as a wake-lag measured in degrees along the toroidal coordinate.

By varying the wake-lag it is possible to control the phase relationship between fuel delivery and the propagating crest. In this work, lag scans are performed across the full angular domain to determine whether the persistence of the crest depends on this phase relationship.

Two reference cases are used for comparison:

- a no injection baseline
- asynchronous injection in which fuel is added without phase alignment

These cases provide a control against which phase-aligned injection can be evaluated.

2.3. Transport and exhaust

Transport along the ring is represented through a combination of directed advection and diffusive smoothing acting on energy and fuel density gradients, consistent with standard advection-diffusion representations of transport processes [13].

Particle exhaust is implemented through simple removal terms with characteristic timescales for fuel and helium ash. These represent the loss of particles from the system and prevent unbounded accumulation during long simulations.

2.4. Operational Parameter Set

Unless otherwise stated, the simulations reported here use the following baseline parameter set. These values were selected to provide stable crest formation while maintaining clear propagation dynamics.

Group	Parameter / rule	Baseline value / setting	Notes
Domain / resolution	Primary production grid	N = 1440	Main resolution for lag scans
Domain / resolution	Ring geometry	1D periodic ring	Reduced conceptual model
Crest geometry	Crest width	0.1	Baseline value used in Monte Carlo reference set
Crest geometry	Crest amplitude	25.0	Baseline value used in Monte Carlo reference set
Transport	Diffusion coefficient	0.01	Baseline value used in Monte Carlo reference set
Transport	Advection scheme	MUSCL-Hancock + minmod TVD	Resolution-consistent transport implementation
Fusion / heating	Fusion reactivity	Bosch-Hale DT fit	Used for $\langle\sigma v\rangle_{DT}(T)$
Fusion / heating	Fusion gain coefficient	G = 108	Baseline operational value
Fusion / heating	Alpha fraction	α=0.2	Forward-biased redistribution kernel
Fuel model	Injection temperature mode	Fixed	Baseline uses fixed-temperature injection
Fuel model	Fuel injection temperature	Tfuel=0.0	Baseline operational value
Fuel model	Thermalisation efficiency	$\epsilon_{\mathrm{therm}} = 1.0$	Thermalised fuel advects at $\epsilon_{\mathrm{therm}} v_{\mathrm{crest}}$
Fuel model	Thermalisation rate	λtherm=5.0	First-order fresh-to-thermalised transfer
Fuel model	Fuel boost density	0.05	Baseline value used in Monte Carlo reference set
Fuel model	Injection width	2.0∘	Baseline value used in Monte Carlo reference set
Injection logic	Injector layout	Fixed injectors, uniformly spaced	Injector density scales with resolution
Injection logic	Wake target	ϕtarget=ϕcrest +Δϕlag	Positive lag = behind crest
Injection logic	Wake acceptance window	max(1∘,Δϕcell)	Criterion for wake-mode firing
Injection logic	Per-orbit gating	Max one fire per injector per orbit	Used in both wake and asynchronous modes
Injection logic	Deposition shape	Normalized Gaussian plume	Conserves injected mass and enthalpy
Exhaust	Exhaust enabled	Yes	Baseline operational setting
Exhaust	DT pumping timescale	τpump,DT =10.0	Baseline operational value
Exhaust	He pumping timescale	τpump,He=5.0	Baseline operational value
Diagnostics / stopping	Stall threshold	0.5 rad per model time unit	Collapse criterion
Diagnostics / stopping	Minimum peak temperature	10.5 keV	Collapse criterion
Diagnostics / stopping	Flatness threshold	1.02	Peak-to-mean temperature ratio criterion

Table 1. Full baseline implementation summary for the reduced propagation model. Baseline values and implementation settings used for the principal wake-lag scans and control runs unless otherwise stated. The table combines the operational parameter set introduced in Section 2.4 with the explicit numerical, injection, and stopping-rule definitions recorded in Appendix C. Most production runs were performed at a grid resolution of $N = 1440$, which was adopted as the primary resolution for the reported lag scans and comparison cases.

Although temperature is reported in keV and DT reactivity is evaluated using the Bosch-Hale fit, the present system remains a reduced mechanism

model in arbitrary units. Parameters such as the fusion-gain coefficient, thermalisation and pumping rates, and collapse thresholds are closure choices introduced to define and interrogate the propagation dynamics of the model rather than calibrated reactor-operating values. The quantitative thresholds and lifetimes reported below should therefore be interpreted as model-specific outcomes, not as predictive requirements for physical plasma devices.

2.5. Numerical resolution

The toroidal domain is discretized with grid sizes ranging from $N = 360$ to $N = 2880$ cells. Convergence tests indicate that the crest structure is effectively resolved by $N = 1440$, and the qualitative wake-lag behaviour remains stable across the tested higher resolutions. Most parameter scans therefore use $N = 1440$ as the primary resolution, with $N = 720$ used for exploratory runs and $N = 2880$ for higher resolution verification.

3. Results

The behaviour of the reduced propagation model is examined through a series of numerical experiments. The primary observable is the persistence of the travelling temperature crest, measured by the runtime of the simulation before the crest loses coherence according to the termination criteria described in Section 2.

The results are presented in stages. First, baseline behaviour is established using simulations with no fuel injection and with asynchronous injection that does not track the crest position. The effect of phase-aligned fuel injection is then examined through a wake-lag scan in which the angular offset between the injection location and the crest is varied across the toroidal domain. Additional tests explore parameter sensitivity, robustness under moderate perturbations of model parameters, and numerical convergence with grid refinement.

3.1. Baseline Behaviour Without Phase Alignment

Baseline behaviour was first examined using two control cases: a simulation with no fuel injection and a simulation with asynchronous injection that does not track the crest position. Both runs were performed at a resolution of $N = 1440$ cells using the operational parameter set described in Section 2.

The crest temperature evolves similarly in both injection cases (Figure 1). In both simulations the initial temperature peak forms rapidly and then decays gradually as the system evolves. The two curves follow nearly identical trajectories throughout the simulation.

Both runs terminate at $t = 26.5$ model time units when the crest loses coherence according to the flatness criterion defined in the model. Asynchronous injection does not extend the lifetime of the propagating structure and collapses on essentially the same timescale as the no injection baseline, even though it receives the same scheduled fuel input per orbit as the phase-aligned cases, with only minor local variations arising from thermalisation and exhaust effects. These baseline results establish that fuel delivery alone does not sustain propagation in the reduced system when the injection location is not phase-aligned with the travelling burn structure. The behaviour of the model therefore provides a suitable control against which phase-dependent injection can be evaluated.

Case	Grid (N)	Injection Mode	Runtime (model units)	Notes
Baseline	1440	None	26.5	crest collapse
Control	1440	Asynchronous	26.5	identical fuel input

Table 2. Baseline runtime comparison for the reduced propagation model. Both simulations use a grid resolution of $N = 1440$. The baseline case contains no fuel injection, while the control case introduces fuel asynchronously without phase alignment to the crest position. Despite continuous fuel input in the asynchronous run, the propagation lifetime remains unchanged relative to the no injection case, indicating that fuel delivery alone does not sustain the travelling burn structure.

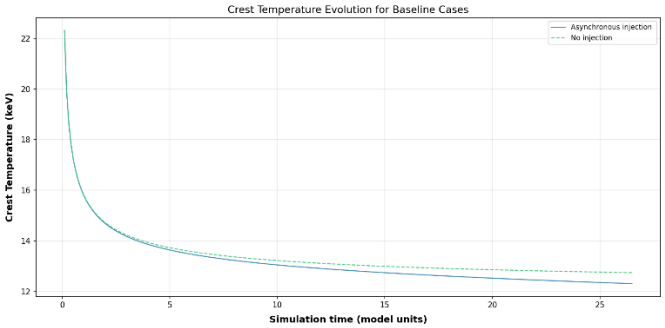

Figure 1. Crest temperature evolution for the baseline simulations. Crest temperature as a function of simulation time for the no injection case and the asynchronous injection control. The two trajectories follow nearly identical decay behaviour, and both terminate at $t = 26.5$ model time units when the crest loses coherence according to the flatness criterion. This confirms that asynchronous fuel input does not extend the lifetime of the propagating structure.

NOTE: High-resolution versions of all figures are provided in Supplement A.

3.2. Phase-Dependent Propagation

The effect of injection timing was examined through a wake-lag scan in which the angular offset between the injection location and the crest position was varied in increments of 0.5°. All runs used the baseline parameter set described in Section 2 with a grid resolution of $N = 1440$. The scan covers the full periodic domain, allowing the behaviour of the propagating structure to be evaluated as a function of injection phase. Under the chosen termination criteria all simulations eventually collapsed, allowing simulation runtime to be used as a consistent metric for comparing phase-dependent propagation behaviour across the lag range.

Burn behaviour varies strongly with injection phase. Injection near the crest produces the highest final burn conditions, with crest temperatures reaching approximately 12.58 keV and reduced fuel consumption per orbit. Despite this strong local burn, the propagating structure collapses quickly under the flatness criterion, producing short runtimes of roughly 10 - 30 model time units across the region extending approximately from -60° to +60°. These behaviours are summarised in Figure 2.

Outside this region the behaviour changes abruptly. For the baseline termination criteria, the transition on the positive-lag branch is confined to the narrow interval 63.5°–64.5°, placing it near 64° at the present scan resolution. Runtime increases rapidly across this interval, rising from 66 model time units at 63.5° to 160.5 model time units at 64.5°. This marks the onset, in the baseline configuration, of a delayed-collapse regime in which the crest remains coherent for significantly longer periods before the 10.5 keV temperature termination condition is reached, rather than the flatness criterion. The transition occurs near 64° in the baseline scan, marking the entry into the longer-lived regime at the present scan resolution. A similar transition appears on the negative-lag branch between about -63° and -64°, placing the corresponding boundary near -63.5° in the baseline scan. The two transitions are broadly symmetric but not identical; runtimes on the positive side of the scan are slightly longer than those observed at the corresponding negative offsets.

The longest runtimes occur near the periodic boundary of the lag domain. The global maximum is found at -179.5° lag, which corresponds to the periodic equivalent (+180.5°). A nearly identical value occurs at -178.5° (+181.5°), with both runs reaching 179 model time units before the crest fails the temperature termination criterion. These locations lie slightly inside the anti-crest position of 180°, suggesting that, within the present closure and scan resolution, the longest runtimes occur just offset from the purely opposite position around the ring.

The scan therefore reveals three distinct regimes: a short-lived propagation region near crest alignment, a sharp transition near ±64° lag, and a long-lived regime extending toward the periodic boundary of the domain. Although all runs eventually terminate due to either temperature or crest flattening, the lifetime of the propagating structure varies by more than an order of magnitude depending on the phase relationship between fuel injection and the travelling burn crest.

Key numerical values from the scan are summarised in Table 3, while Figure 2 shows the full lag scan and Figure 3 provides a zoomed view of the transition region.

Wake-Lag (deg)	Periodic Equivalent (deg)	Final Crest Temp (keV)	Runtime (model units)	Fuel / Orbit	Notes
−10.0	+350.0	12.55	11.5	0.588	near crest region
−5.0	+355.0	12.56	11.0	0.596	strong burn, short lifetime
0.0	0.0	12.58	12.0	0.505	peak crest temperature
+5.0	+5.0	12.55	11.5	0.575	near crest region
+10.0	+10.0	12.56	11.5	0.580	near crest region
−63.0	+297.0	11.90	65.0	0.561	pre-transition
−63.5	+296.5	11.88	70.0	0.558	transition region
−64.0	+296.0	11.75	92.0	0.553	transition onset
+63.5	+63.5	11.88	66.0	0.560	pre-transition
+64.0	+64.0	11.70	92.0	0.552	transition onset
+64.5	+64.5	10.50	160.5	0.544	long-lived regime begins
−179.5	+180.5	10.50	179.0	0.545	maximum runtime
−178.5	+181.5	10.50	179.0	0.545	periodic equivalent peak

Table 3. Representative values from the wake-lag parameter scan. Injection close to crest alignment produces the highest final crest temperatures and lowest fuel cost per orbit but results in rapid collapse of the propagating structure. A sharp transition occurs near ±64° lag where runtime increases rapidly. The longest propagation lifetimes occur near the periodic boundary of the domain close to ±180°, where runtimes approach 179 model time units before termination.

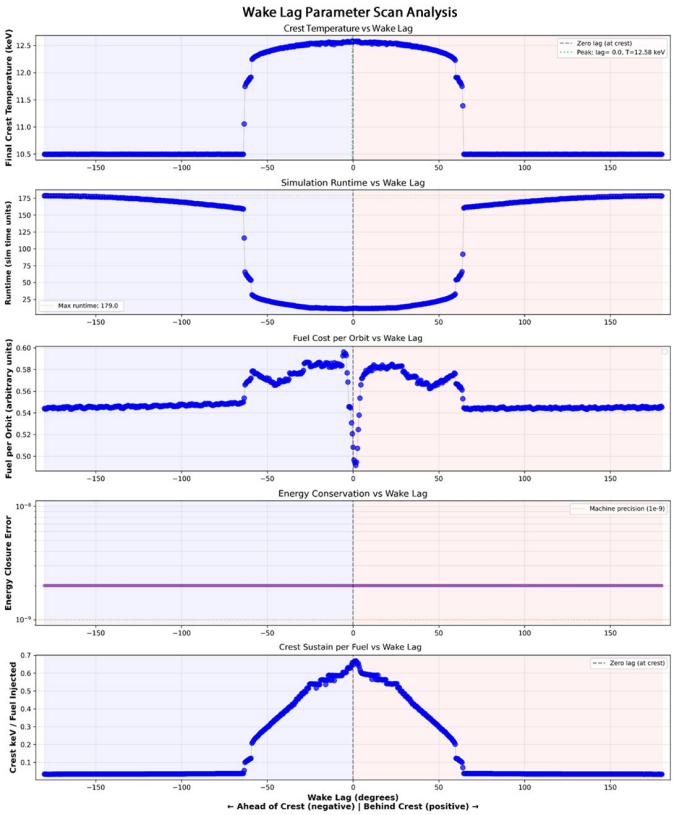

Figure 2. Wake-lag parameter scan for phase-dependent fuel injection (N = 1440). Each point corresponds to a simulation with constant fuel input but different injection phase relative to the travelling burn crest. Negative lag denotes injection ahead of the crest and positive lag denotes injection behind it. Panels show, from top to bottom: final crest temperature, runtime before termination, fuel cost per orbit, energy conservation error, and crest sustainment per unit fuel injected.

Injection near crest alignment (lag ≈ 0°) produces the highest crest temperatures but short runtimes, indicating that fuelling alone does not stabilise the propagating structure. A sharp transition occurs near |lag| ≈ 64°, beyond which runtimes increase rapidly. The longest lifetimes occur near the anti-crest region close to ±180°, where runtimes approach ~179 model time units before termination. Energy closure remains stable at approximately 10^{-9} across the scan, indicating numerical stability of the simulations.

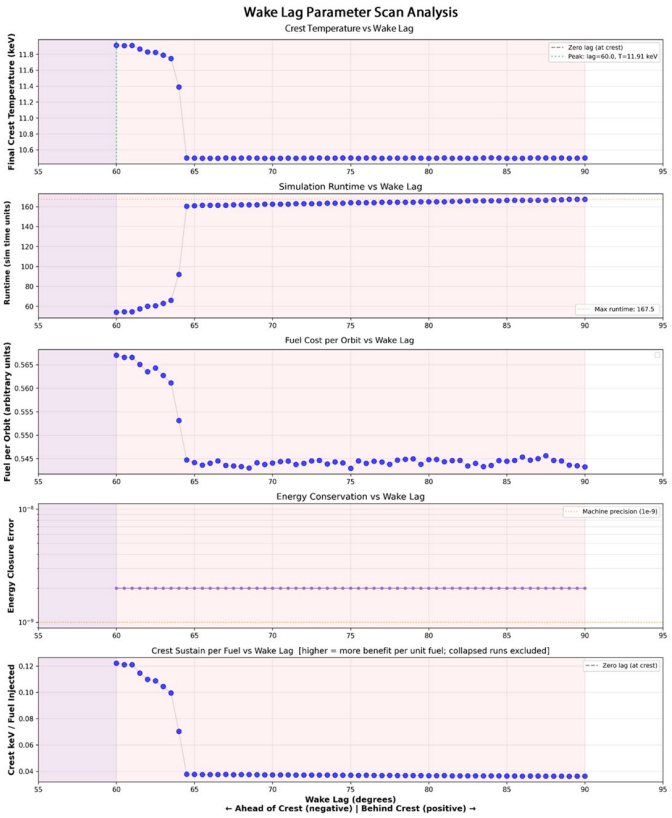

Figure 3. Zoomed view of the wake-lag transition region near |lag| ≈ 64° (N = 1440). The figure shows a detailed view of the parameter scan in the region where the system transitions between short-lived crest-aligned behaviour and long-lived anti-crest propagation. A sharp transition occurs between approximately 63.5° and 64.5° lag. For lags below this boundary the crest temperature remains elevated (≈11.8 - 11.9 keV) and runtimes remain short (≈55 - 65 model time units). Beyond the transition the crest temperature becomes the termination criterion (≈10.5 keV) not the flatness threshold while runtimes increase rapidly to ≈160 - 170 model time units.

The transition is narrow (≈1° wide) and appears on both positive and negative lag branches with slight asymmetry, the positive branch exhibits marginally longer runtimes than the negative branch. Fuel consumption per orbit decreases after the transition as the propagating structure stabilises, while energy closure remains constant at $\approx 10^{-9}$, indicating numerical stability of the scan.

3.3. Crest Structure and Burn Behaviour

The extended burn lifetimes observed in the wake-aligned regime arise from a competition between two opposing effects of cold fuel injection on the propagating crest. When fuel is deposited close to the crest maximum, injected mass enters the region where the local temperature gradient is steepest, and the crest structure is most sensitive to perturbation. Although fusion heating is locally enhanced, the thermalisation timescale of the cold fuel is not short relative to the crest transit time across the injection region. The net effect is dilution of the reacting population faster than the crest can thermalise the incoming fuel, and the propagating structure collapses despite a transient increase in local temperature, consistent with the delayed thermalisation mechanism introduced in Section 2. This quench mechanism explains why injection near zero lag produces the highest final crest temperatures but the shortest propagation lifetimes.

Beyond offsets of about ±64° in the baseline scan, the crest has already transited the injection location. The wake retains sufficient temperature to thermalise incoming fuel on a timescale short relative to the orbital period (the time required for the crest to complete one toroidal circuit), such that injected fuel is heated into the reacting population before the crest returns. Each successive crest transit therefore encounters pre-conditioned fuel in its wake rather than cold mass in its core. This repeated entrainment sustains coherent phase-organised propagation over much longer durations and produces the persistence plateau observed in the lag scans. In the baseline closure, the observed transition near ±64° is consistent with a thermal-entrainment threshold: an angular separation at which cold-fuel injection appears to shift from a crest-quench mechanism to a wake-entrainment mechanism. Consistent with the lag-scan diagnostics, this boundary also marks a change in the dominant termination mode: below it, runs typically end through loss of crest coherence under the flatness criterion, whereas beyond it coherent propagation persists long enough for the temperature threshold to become limiting.

To examine this mechanism quantitatively, the internal dynamics of the propagating crest were analysed for three representative regimes: a collapsing regime (wake-lag = 50°), a near-threshold regime just below the transition (60°), and a sustained regime well beyond the transition (70°).

Crest propagation behaviour differs systematically across these regimes. The temporal evolution of crest angular velocity is shown in Figure 4. In the collapsing case the crest progressively decelerates relative to the nominal propagation rate, with a measured drift of -5.7×10^{-3} radians per model time unit^{-2} (based on model-internal diagnostics). The transitional regime shows reduced drift (-2.6×10^{-3} radians per model time unit $^{-2}$), while

the sustained regime exhibits near-zero drift (-3×10^{-4} radians per model time unit $^{-2}$). This systematic reduction in velocity drift directly reflects the transition from quench-dominated to entrainment-dominated injection. When cold fuel is deposited into the crest core the additional mass loading decelerates the front; when the same fuel is deposited into the wake it is pre-thermalised before the next transit and the crest propagates at near-nominal speed.

The persistence of the burn front is reflected in the temperature contrast between the crest peak and the background plasma (Figure 5). In all regimes the initial contrast decreases rapidly due to thermalisation. However, the decay rate differs substantially across regimes. In the collapsing case the contrast quickly approaches the flatness threshold that terminates the run, consistent with ongoing dilution of the crest core by cold injected mass. In the sustained regime a small but persistent temperature gradient is maintained for a significantly longer duration, indicating that the burn front remains spatially coherent and that entrainment of wake fuel is continuously reinforcing the crest against thermal diffusion.

Fuel injection statistics were examined to confirm that these lifetime differences are not driven by changes in fuelling rate. The distribution of inter-fire intervals for the three regimes is shown in Figure 6. The mean cadence remains close to one orbit in all cases (1.053, 1.050, and 1.097 orbits respectively) with comparable variance. The extended burn lifetime therefore reflects a change in the effectiveness of each injection event rather than any increase in fuel quantity or delivery frequency.

The entrainment mechanism is further constrained by the phase relationship between the propagating crest and the injection forcing. Across all three representative cases, the residual phase offset remains approximately constant over the duration of each run, indicating that the crest maintains a nearly fixed phase relation to the forcing clock while coherent propagation persists (Figure 7). The principal difference between regimes is therefore not the existence of an instantaneous phase relationship, but its durability: in the collapsing cases the crest structure decays rapidly despite this near-constant offset, whereas in the sustained regime the same phase organisation is maintained over a much longer interval. The transition from quench to entrainment is thus better interpreted as a transition from short-lived to persistent phase-organised propagation.

Taken together these diagnostics describe a single coherent picture. In the baseline configuration, a threshold near ±64° separates two dynamical regimes distinguished not by how much fuel is delivered but by where relative to the crest that fuel is thermalised. Inside the threshold, cold mass loads the crest core, velocity drift increases, and temperature contrast collapses before long-lived propagation can be sustained. Outside the threshold, the wake pre-conditions the injected fuel, drift approaches zero, a persistent temperature gradient is maintained, and the burn front sustains a durable phase-organised propagation state with the injection forcing.

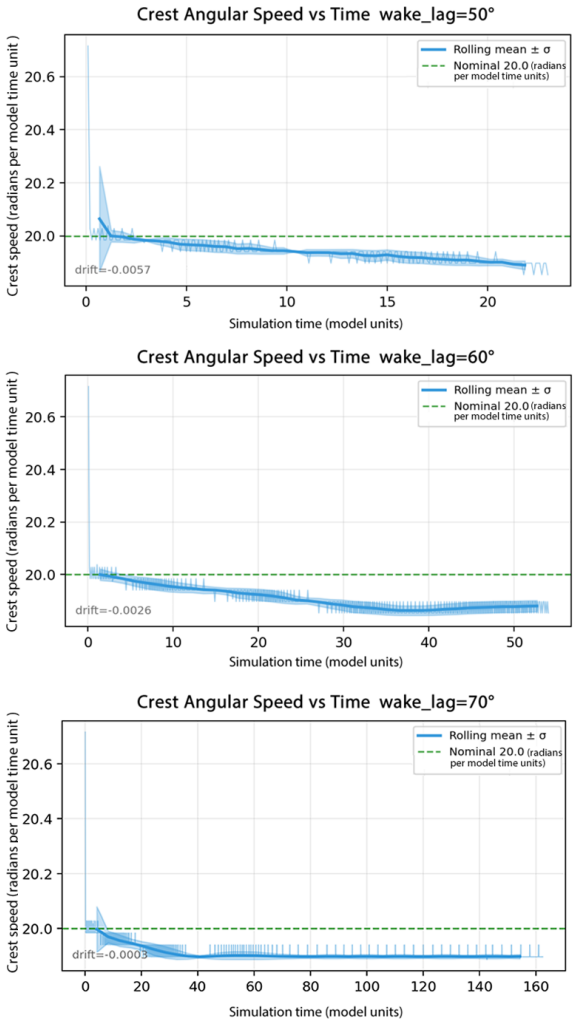

Figure 4. Crest angular speed evolution for representative wake-lag regimes. Temporal evolution of the burn crest angular velocity for three wake-lag values representing collapsing (50°), near-threshold (60°), and sustained (70°) propagation regimes. The solid line shows the rolling mean crest speed with ±σ envelope, while the dashed line indicates the nominal propagation speed (20 radians per model time unit^{-1}). As the wake-lag approaches and passes the thermal entrainment threshold, the drift of the crest relative to the nominal speed decreases significantly (-5.7×10^{-3}, -2.6×10^{-3}, and -3×10^{-4} radians per model time unit^{-2} respectively), indicating progressively more dynamically stabilised propagation.

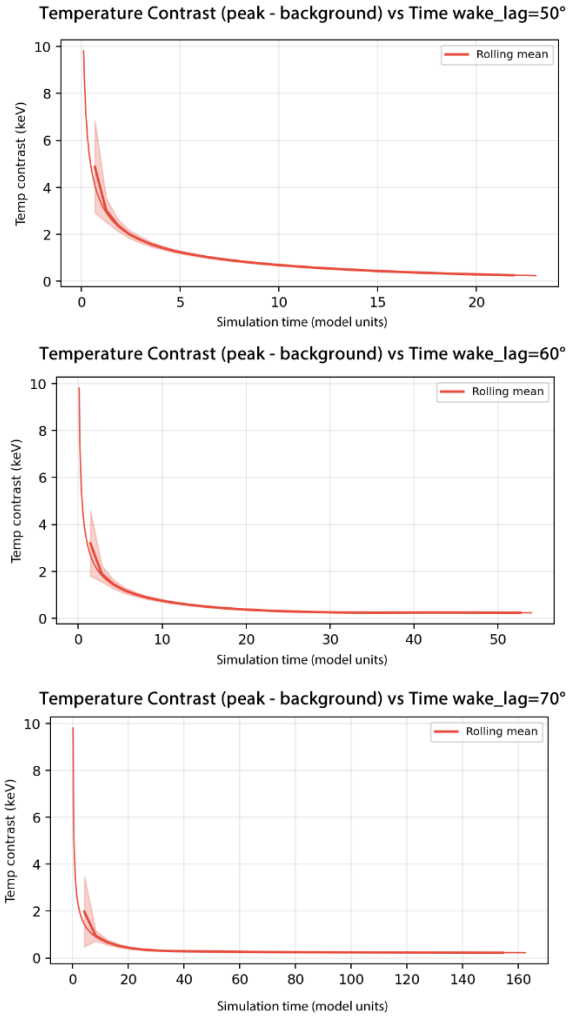

Figure 5. Evolution of crest temperature contrast for representative wake-lag regimes. Temperature contrast between the burn crest peak and the background plasma as a function of simulation time for wake-lags of 50°, 60°, and 70°. The solid line shows the rolling mean contrast with early-time variability indicated by the shaded region. In all cases the contrast decreases due to thermalisation; however, the sustained regime (70°) maintains a small but persistent gradient for a significantly longer duration than the collapsing and near-threshold cases, indicating prolonged spatial coherence of the propagating burn front.

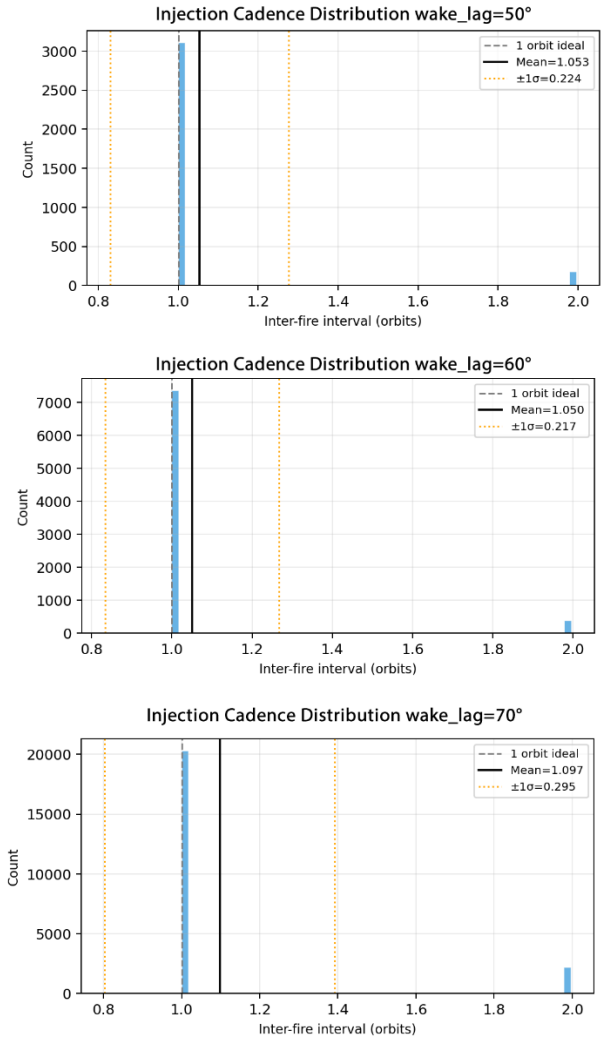

Figure 6. Distribution of injection cadence for representative wake-lag regimes. Histogram of inter-fire intervals (in orbits) for wake-lags of 50°, 60°, and 70°. The dashed line indicates the ideal cadence of one orbit, while the solid line marks the measured mean interval and the dotted lines show ±1σ. The mean cadence remains close to one orbit in all regimes (1.053, 1.050, and 1.097 orbits respectively), indicating that the observed differences in burn lifetime are not driven by changes in fuel injection rate.

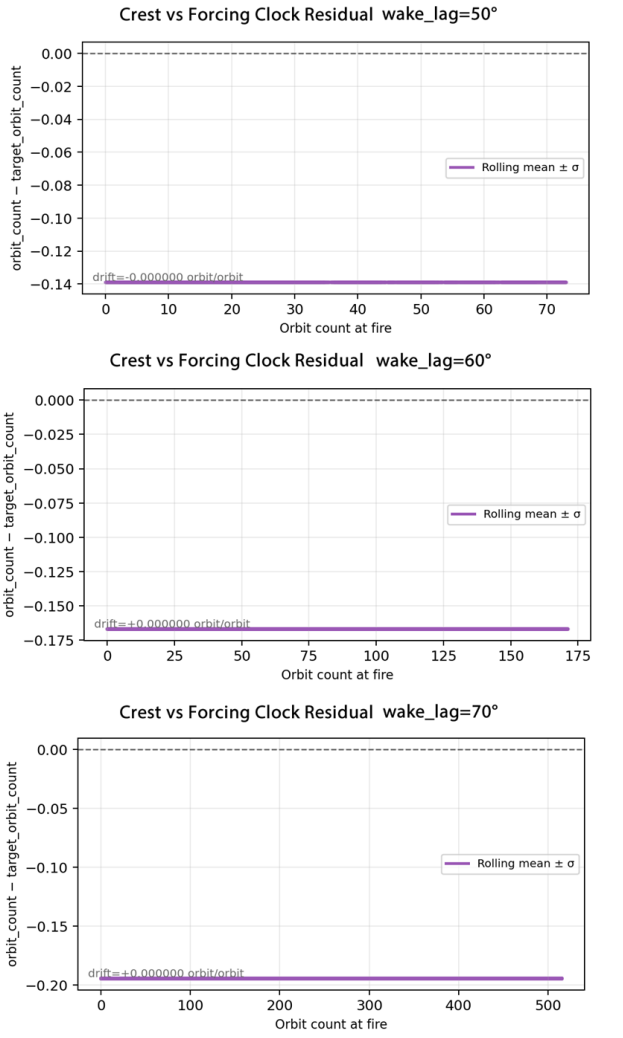

Figure 7. Residual between crest orbit count and forcing-clock expectation. Residual between the actual crest orbit count and the target orbit count prescribed by the forcing schedule at each injection fire event for wake-lags of 50°, 60°, and 70°. The solid line shows the rolling mean residual with ±σ envelope, and the dashed line indicates zero residual. In each case the residual remains approximately constant with negligible drift over the observed lifetime of the run, indicating that the propagating crest maintains a nearly fixed timing offset relative to the injection forcing while coherent propagation persists.

3.4. Parameter Sensitivity

To evaluate whether the observed behaviour depends on a narrow choice of model parameters, several sensitivity tests were performed around the baseline configuration described in Section 2.4. These tests examine how variations in selected parameters affect crest persistence and the structure of the wake-lag response. The aim is not to exhaustively explore the parameter space, but to verify that the phase-dependent behaviour identified in the lag scans remains qualitatively stable under moderate perturbations of the model assumptions. Parameters examined include the simulation termination criteria and selected dynamical parameters governing crest propagation and fuel coupling. Results are presented in the following subsections.

3.4.1. Termination Criteria Sensitivity

The persistence of the propagating crest in the reduced model is determined by termination criteria that detect loss of coherent structure. Because the measured simulation lifetime depends on these stopping rules, it is important to verify that the wake-aligned persistence behaviour does not arise from a particular choice of termination definition.

Sensitivity tests were therefore performed in which each termination criterion was varied independently while all other parameters were held at their baseline values. The tested ranges represent moderate perturbations around the baseline thresholds used in the primary simulations. Specifically, the following parameters were varied:

- crest flatness threshold
- minimum crest propagation speed
- minimum crest temperature

The values tested are summarized in Table 4.

Criterion	Values Tested
Crest flatness threshold	1.01, 1.02, 1.03
Minimum crest propagation speed	0.3, 0.5, 0.7
Minimum crest temperature (keV)	10.3, 10.5, 10.7

Table 4. Termination criteria sensitivity tests used to evaluate the robustness of crest persistence measurements. Each criterion was varied independently around the baseline configuration while all other parameters were held fixed. The values listed correspond to the thresholds used in the wake-lag scans shown in Figure 8.

Across all tested cases, the overall qualitative structure of the wake-lag response remains intact. In every configuration the lag scans exhibit three characteristic regions: a short-lived propagation regime near crest-aligned injection, a sharp transition region, and an extended wake-aligned persistence branch in which the travelling crest survives significantly longer. This indicates that the existence of the phase-dependent persistence plateau does not depend on a particular stopping rule.

The minimum crest temperature and minimum propagation speed thresholds produce only minor shifts in the lag-response curves and leave the location of the transition largely unchanged. By contrast, the crest flatness threshold has a stronger influence on the detected collapse boundary. Increasing the flatness requirement shifts the transition point progressively farther into the wake region, while a more permissive threshold moves the boundary closer to the crest. This behaviour is expected because the flatness criterion directly measures the structural coherence of the propagating crest and therefore most strongly affects when the simulation declares the structure to have collapsed.

Despite these quantitative shifts in the detected transition location, the underlying phase-dependent structure of the wake-lag response remains consistent. In all cases a distinct wake-aligned region of enhanced persistence is present, while injections placed sufficiently close to the crest continue to produce rapid collapse. The sensitivity tests therefore indicate that the extended runtimes observed in the wake-aligned regime arise from the intrinsic crest dynamics of the reduced model rather than from a particular termination definition.

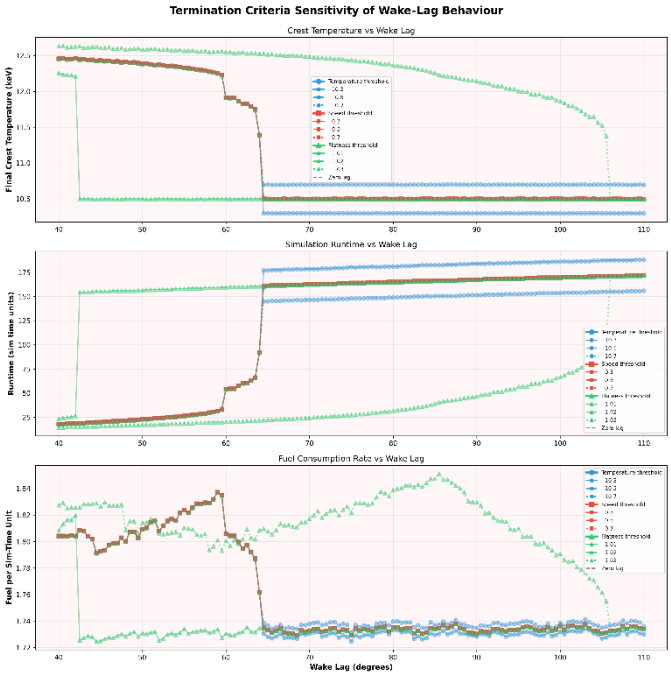

Figure 8. Termination Criteria Sensitivity of Wake-Lag Behaviour. Wake-lag scans repeated while independently varying the termination thresholds used to define crest collapse. Temperature thresholds of 10.3, 10.5, and 10.7 keV; minimum crest speed thresholds of 0.3, 0.5, and 0.7; and crest flatness thresholds of 1.01, 1.02, and 1.03 were tested while all other parameters were held at baseline values. Panels show crest temperature (top), simulation runtime (middle), and fuel consumption rate (bottom) as functions of wake-lag. Although the precise lag at which collapse is detected shifts slightly with the threshold choice, the wake-aligned persistence region and overall lag structure remain consistent across all criteria.

3.4.2. Crest Speed Sensitivity

The propagation speed of the temperature crest determines the rate at which fuel injected into the wake is advected relative to the burn front.

Because wake-aligned fuelling relies on the timing between injection and crest motion, variations in the crest speed could potentially alter the phase window in which reinforcement occurs.

To test the sensitivity of the wake-alignment behaviour to this parameter, the wake-lag scan was repeated while varying the crest speed parameter. Values of 16, 20, and 24 were examined while all other parameters were held fixed at their baseline values.

Across the tested range the overall structure of the wake-lag response remains consistent. A narrow region of enhanced persistence continues to appear behind the crest, while injections ahead of the crest collapse rapidly. Changing the crest speed primarily alters the absolute runtime and the sharpness of the transition to collapse but does not eliminate the persistence plateau associated with wake-aligned injection.

This indicates that the phase-dependent reinforcement observed in the model is not dependent on a single propagation velocity but persists across moderate variations in crest speed.

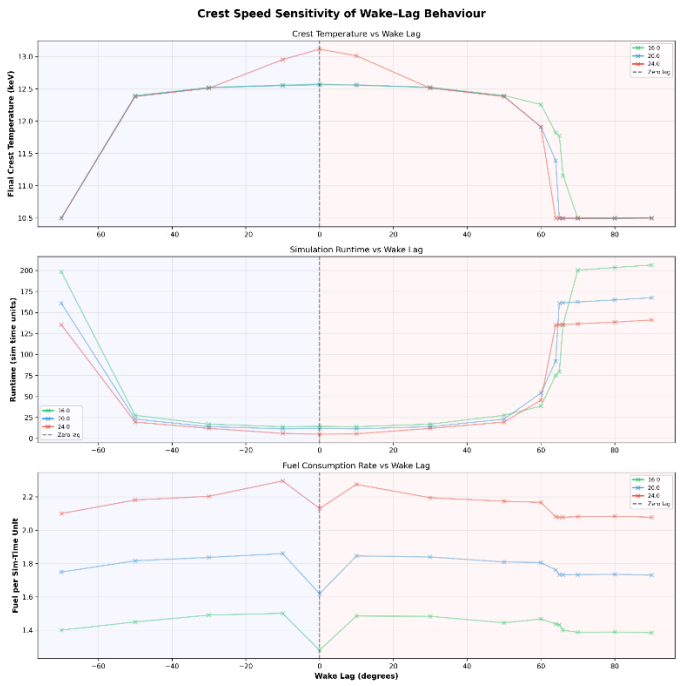

Figure 9. Crest Speed Sensitivity of Wake-Lag Behaviour. Wake-lag scans repeated while varying the crest propagation speed parameter. Crest speed values of 16, 20, and 24 radians per model time unit $^{-1}$ were tested while all other model parameters were held at their baseline values. Panels show crest temperature (top), simulation runtime (middle), and fuel consumption rate (bottom) as functions of wake-lag. Changes in crest speed alter the absolute runtime and fuel consumption levels but do not remove the characteristic wake-aligned persistence region. The overall lag structure remains similar across all tested speeds, indicating that the phase-dependent reinforcement mechanism is not tied to a single propagation velocity.

3.4.3. Thermalisation Rate Sensitivity

The wake-aligned fuelling mechanism in the reduced model depends on a delay between the injection of fresh fuel and its participation in fusion reactions. This delay is controlled by the thermalisation rate parameter, which governs how rapidly injected particles join the reacting population. Variations in this parameter alter the temporal coupling between fuel injection and crest propagation and therefore provide a useful test of whether the observed phase-dependent behaviour depends on a specific delay timescale.

To evaluate this sensitivity, the wake-lag scan was repeated while varying the thermalisation rate parameter around its baseline value. Values of 1, 2, 3, 5, 7.5, 10, 15 and 20 were examined while all other parameters were held fixed.

Across this range the qualitative structure of the wake-lag response remains present. A region of enhanced persistence continues to appear behind the crest, while injections ahead of the crest collapse rapidly. Changing the thermalisation rate primarily alters the sharpness of the persistence plateau and the absolute runtime achieved within the wake-aligned region. Faster thermalisation reduces the delay between injection and fusion participation, slightly narrowing the region of effective wake capture, while slower thermalisation allows injected fuel to remain in the wake for longer before contributing to the burn.

Despite these shifts, the existence of a wake-aligned persistence region remains robust across the tested range. This indicates that the phase-dependent reinforcement mechanism does not depend on a single tuned thermalisation timescale but persists across moderate variations in the delay between injection and fuel activation.

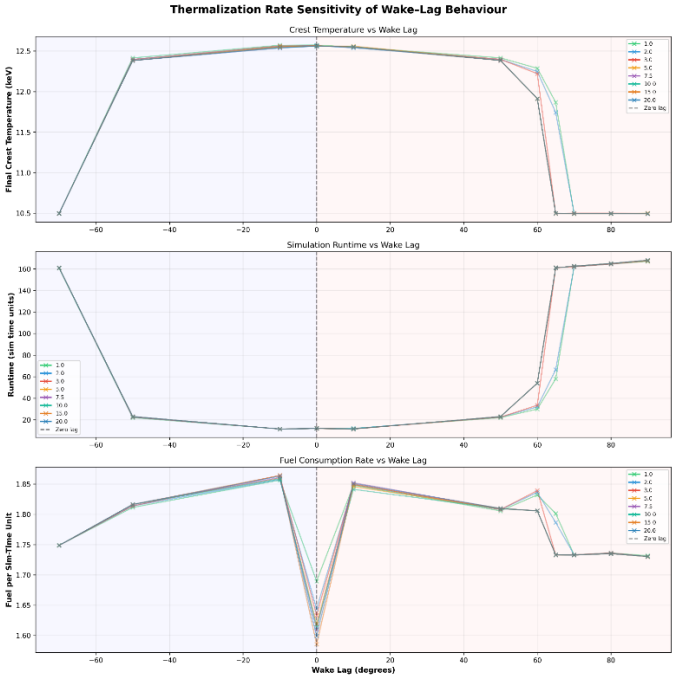

Figure 10. Thermalisation Rate Sensitivity of Wake-Lag Behaviour. Wake-lag scans repeated while varying the thermalisation rate parameter controlling the delay between fuel injection and fusion participation. Thermalisation rates of 1, 2, 3, 5, 7.5, 10, 15 and 20 were tested while all other parameters were held at their baseline values. Panels show crest temperature (top), simulation runtime (middle), and fuel consumption rate (bottom) as functions of wake-lag. Although the width and sharpness of the persistence region vary slightly with the thermalisation rate, the wake-aligned plateau remains present across the tested range, indicating that the phase-dependent reinforcement mechanism is robust to moderate variations in the injection-fusion delay timescale.

3.4.4. Monte Carlo Ensemble Robustness

To probe the robustness of the wake-lag response more broadly than the single-parameter sweeps above, two Monte Carlo ensembles were generated around the baseline configuration using a fixed random seed for reproducibility. For each sampled parameter set, the wake-lag scan was repeated at a reduced set of representative offsets spanning the principal regions of interest: −70°, −50°, −30°, −10°, 0°, 10°, 30°, 50°, 60°, 65°, 70°, 80°, and 90°. All runs were performed at the primary production resolution using the same baseline configuration as elsewhere in this study. In the reported Monte Carlo ensembles, the thermalised-fuel entrainment factor was held fixed at its baseline value, $\epsilon_{therm}=1.0$, and was not included among the perturbed parameters.

In Pass A, parameters associated with thermalisation and pumping kinetics were perturbed simultaneously by ±10% about baseline. These were the thermalisation rate and the deuterium-tritium and helium pumping timescales. In Pass B, parameters associated with crest geometry and transport were perturbed simultaneously, comprising the crest width, crest amplitude, fuel boost density, injection width, and diffusion. A smaller perturbation range of ±5% was used for Pass B in order to test local robustness of these structural and transport terms without substantially altering the imposed crest shape itself.

The Monte Carlo results preserve the same broad conclusion as the deterministic sensitivity tests: the response remains phase-dependent under simultaneous moderate perturbations of multiple model parameters. Across both ensembles, the curves remain structured as functions of wake-lag rather than collapsing into uncorrelated scatter. Final crest temperature, runtime, and fuel consumption continue to vary systematically with injection phase, indicating that the lag dependence is not an artefact of a single finely tuned parameter choice.

The two ensembles do, however, differ modestly in the degree of spread they introduce. In the Pass A ensemble, the sampled curves remain tightly clustered across most of the tested lag range. The qualitative wake-lag structure is preserved almost unchanged from sample to sample, with the principal variation concentrated near the positive-lag transition around +60° to +65°. These results indicate that moderate simultaneous perturbations to the thermalisation and pumping kinetics do not remove the phase ordering of the response and only weakly affect the sharpness and exact location of the collapse boundary.

Pass B shows a broader ensemble spread than Pass A, particularly in fuel consumption level and in the sharpness of the positive-lag transition. Even so, the same overall wake-lag structure remains clearly visible across the ensemble: a sustained plateau over the wake-aligned region, rapid collapse for sufficiently positive lag, and the same broad phase ordering seen in the deterministic sweeps. This indicates that, within the perturbation ranges tested, moderate variations in crest-geometry and transport parameters alter quantitative details of the response without

eliminating the underlying phase-dependent reinforcement behaviour.

Taken together, the Monte Carlo ensembles indicate that the phase-dependent wake-lag response is not a fragile consequence of a single baseline parameter set. Moderate simultaneous perturbations preserve the same qualitative dependence on injection phase in both parameter groups. In the ensemble runs shown here, Pass A remains more tightly clustered than Pass B, while both retain the characteristic lag-dependent structure of the wake-aligned regime.

Pass	Parameter group	Parameters varied	Baseline values	Ensemble Size	Perturbation range
A	Thermalisation & pumping kinetics	thermalisation rate; DT pumping timescale; He pumping timescale	5.0; 10.0; 5.0	25	±10%
B	Crest geometry & transport	crest width; crest amplitude; fuel boost density; injection width; diffusion	0.1; 25.0; 0.05; 2.0; 0.01	25	±5%

Table 5. Baseline parameter groups and perturbation ranges used for the Monte Carlo ensemble robustness study. Pass A sampled thermalisation and pumping-kinetics parameters within ±10% of baseline, whereas Pass B sampled crest-geometry and transport parameters within ±5% of baseline. Each sampled parameter set was evaluated over the reduced wake-lag set −70°, −50°, −30°, −10°, 0°, 10°, 30°, 50°, 60°, 65°, 70°, 80°, and 90°.

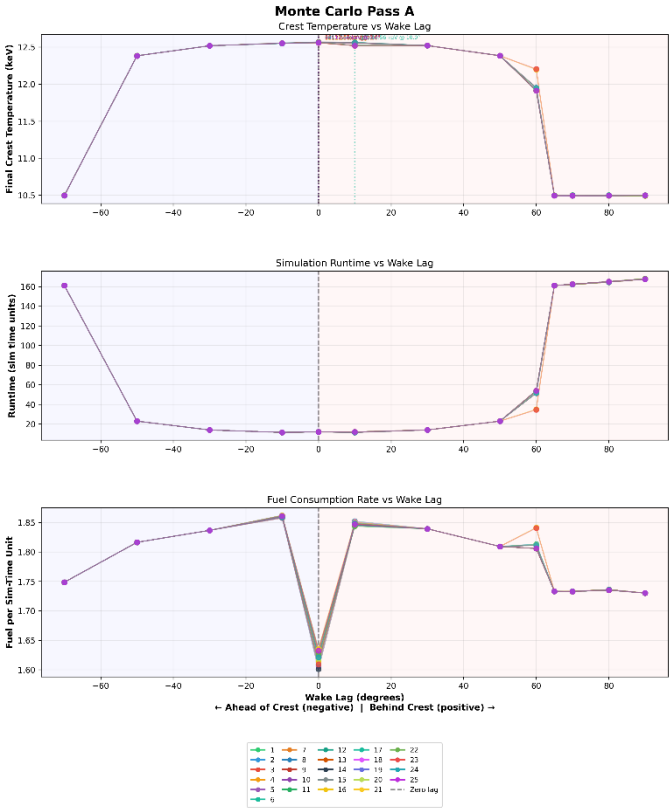

Figure 11. Monte Carlo robustness test, Pass A. Wake-lag scans repeated for an ensemble of randomly perturbed parameter sets in which the thermalisation rate and the deuterium-tritium and helium pumping timescales were varied simultaneously around the baseline configuration. Panels show final crest temperature (top), simulation runtime (middle), and fuel consumption rate (bottom) as functions of wake-lag. The ensemble remains tightly clustered across most sampled lags, with only modest variation appearing near the positive-lag transition around +60° to +65°. The qualitative wake-lag structure is preserved throughout, indicating that the phase-dependent response is robust to moderate simultaneous perturbations of thermalisation and pumping kinetics.

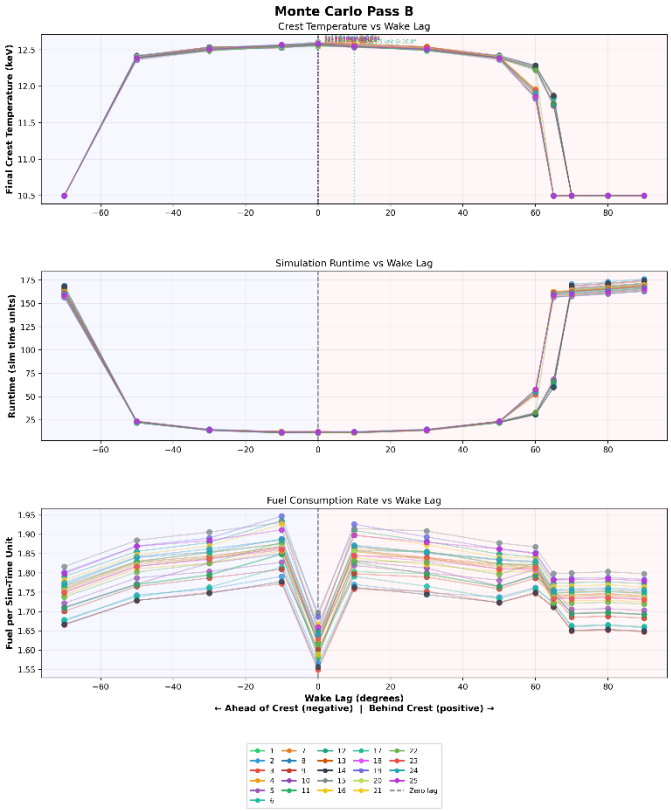

Figure 12. Monte Carlo robustness test, Pass B. Wake-lag scans repeated for an ensemble of randomly perturbed parameter sets in which crest width, crest amplitude, fuel boost density, injection width, and diffusion were varied simultaneously around the baseline configuration. Panels show final crest temperature (top), simulation runtime (middle), and fuel consumption rate (bottom) as functions of wake-lag. The ensemble shows a broader but still well-structured response than Pass A, with the main variation concentrated near the positive-lag transition and in the overall fuel-consumption level. The characteristic wake-lag plateau and collapse structure remain clearly preserved, indicating robustness to moderate simultaneous perturbations of crest-shape and transport parameters.

3.5. Numerical Robustness

Because the reduced propagation model relies on spatial advection and localised crest tracking, it is important to verify that the wake-aligned behaviour is not an artefact of grid resolution. Numerical robustness was therefore examined in two ways: first by repeating the wake-lag scan across multiple grid resolutions, and second by directly comparing the resolved crest structure in the vicinity of the burn peak.

Wake-lag scans were repeated at N = 720, 1440, and 2880 cells while all other parameters were held fixed. The resulting lag-response curves are shown in Figure 13. Across these resolutions the overall structure of the wake-lag response remains consistent. In each case a broad wake-aligned region of enhanced persistence is observed, while injections ahead of the crest collapse rapidly. The collapse boundaries and plateau structure remain similar across the tested resolutions, indicating that the phase-dependent behaviour is not a discretization artefact.

To examine convergence of the local crest structure, the temperature field in the vicinity of the crest was also compared across N = 360, 720, 1440, and 2880 at a fixed simulation time, t = 10 (Figure 14), corresponding to the start of the analysis window used for the representative convergence metrics in Table 6. The N = 360 case is visibly under-resolved and exhibits a broadened, distorted crest profile. The N = 720 case captures the overall crest morphology but remains distinguishable from the higher-resolution cases. By contrast, the N = 1440 and N = 2880 profiles differ only modestly in crest-region shape, indicating that the local crest structure is effectively resolved by N = 1440 for the purposes of the reported wake-lag comparisons.

Quantitative convergence metrics for a representative sustained wake-aligned case are summarized in Table 6. For consistency with the wake-lag diagnostics reported earlier, the final crest temperature used here is the same crest-local metric used elsewhere in the manuscript, defined as the mean temperature within a fixed ±3° window about the hottest cell. Final crest amplitude shows modest upward drift between N = 720, N = 1440, and N = 2880, while the efficiency metrics remain of comparable magnitude across the higher-resolution cases, although they are not strictly monotonic with refinement. Energy closure remains at approximately 10^{-9} in all cases, confirming that the refinement study is not contaminated by numerical instability or conservation drift. These residuals demonstrate that numerical errors in energy and fuel tracking are negligible, ensuring the observed persistence behaviour arises from the model physics rather than numerical artefacts.

Taken together, these results indicate that the principal wake-aligned persistence behaviour reported in this study is numerically robust and that the primary production resolution of N = 1440 is sufficiently resolved for the reported wake-lag comparisons in the reduced model.

Grid Resolution (N)	Final Crest Temperature (keV)	Integrated Efficiency	Burn-Time Efficiency	Closure Error
360	10.625	4.80710	0.41988	2.00e-09
720	11.649	5.33034	0.44101	2.00e-09
1440	12.169	5.07854	0.40541	2.00e-09
2880	12.515	5.23589	0.40606	2.00e-09

Table 6. Numerical convergence metrics across grid resolution for a representative sustained wake-aligned case (lag = 70°; analysis window t = 10-40). Final crest temperature, integrated efficiency, burn-time efficiency, and closure error are reported for simulations performed at resolutions from $N = 360$ to $N = 2880$. Here, final crest temperature denotes the final value of the crest-local temperature diagnostic, defined as the mean temperature within a fixed $\pm 3°$window about the hottest cell. Integrated efficiency is the time integral of this crest-local temperature normalised by total injected fuel, while burn-time efficiency is the duration for which the crest-local peak temperature exceeds 10 keV, also normalised by total injected fuel. In conjunction with Figure 14, the results indicate that the crest structure is effectively resolved by $N = 1440$, while the efficiency measures remain of similar magnitude across the higher-resolution cases.

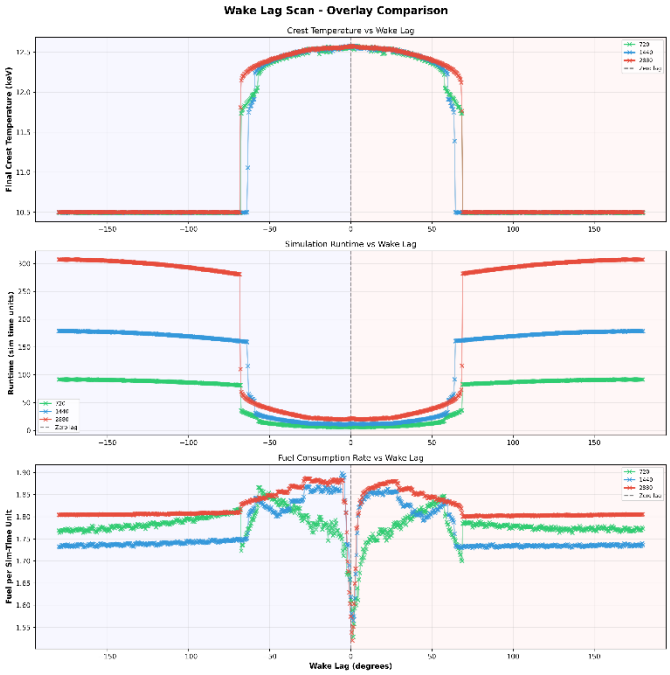

Figure 13. Grid-resolution sensitivity of wake-lag behaviour. Wake-lag scans performed at N = 720, N = 1440, and N = 2880 with all other parameters held fixed. Panels show crest temperature (top), simulation runtime (middle), and fuel consumption rate (bottom) as functions of wake-lag. The overall lag structure remains consistent across resolutions, indicating that the wake-aligned persistence region is not a discretization artefact.

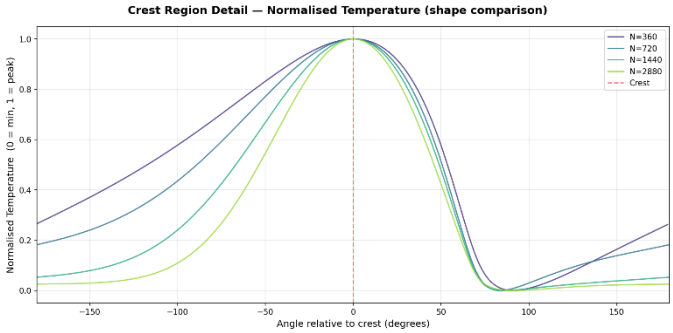

Figure 14. Crest-region detail: normalised temperature profiles (shape comparison at $t = 10$). Temperature profiles in a ±180° window centred on the crest are shown for N = 360, N = 720, N = 1440, and N = 2880 at fixed simulation time. The lowest-resolution case is visibly under-resolved, while the N = 720 case captures the overall crest morphology but remains distinguishable from the higher-resolution profiles. The N = 1440 and N = 2880 cases differ only modestly in crest-region shape, indicating that the crest structure is effectively resolved at N = 1440 for the purposes of the reported wake-lag comparisons.

4. Discussion

The results presented above demonstrate that the persistence of the propagating burn in the reduced model depends strongly on the phase relationship between fuel injection and the moving crest. This behaviour arises from the delayed thermalisation, and transport dynamics defined in the reduced propagation model introduced in Section 2. Across the explored wake-lag scans, a well-defined persistence plateau emerges when injection occurs sufficiently offset from the crest maximum while remaining dynamically coupled to the propagating wave. In contrast, injections placed too close to the crest peak lead to rapid collapse despite identical total fuel input, indicating that direct perturbation of the crest structure shortens propagation lifetime even while locally increasing temperature.

A simple physical interpretation is that the interaction between injection timing, transport, and reaction rate produces either constructive or destructive coupling with the propagating burn structure. When fuel is injected within the wake region of the crest, the local temperature remains elevated, and the fusion reaction rate is already increasing with temperature. Newly injected fuel is therefore rapidly heated and contributes to the propagating front, reinforcing the burn structure. By contrast, injection occurring ahead of the crest introduces fuel into a colder region where diffusion and transport act before significant heating occurs. In this case the fuel disperses

before it can reinforce the propagating front, and the crest subsequently collapses despite identical total fuel input.

Inspection of the spatial temperature fields suggests that the crest itself occupies a relatively narrow angular domain, while the dynamical influence of the travelling wave extends over a broader region that may be interpreted as its effective wake. Injections occurring within this wake-coupled region appear capable of reinforcing the propagating burn front without destabilising the crest core, producing the persistence plateau observed in the lag scans.

In the following sections we interpret this behaviour in terms of phase-coupled propagation dynamics and examine its possible implications for fuelling strategies in systems exhibiting travelling burn structures. Particular attention is given to the geometry of the wake-lag window, the limitations of the reduced model, and the potential relevance of phase-synchronised fuelling concepts in more complex plasma environments.

4.1. Phase Alignment as a Stability Mechanism

The wake-lag scans reveal that localised fuelling events interact with the propagating burn structure in a strongly phase-dependent manner. When injection occurs close to the crest maximum, the local temperature increases but the overall propagation lifetime decreases. This behaviour indicates that direct perturbation of the crest structure can destabilise the propagating wave even when additional fuel is supplied.

By contrast, injections occurring sufficiently offset from the crest peak produce a markedly different outcome. In these cases, the injected fuel enters a region where the temperature is already elevated but the crest itself has not yet been directly perturbed. Because the fusion reaction rate increases rapidly with temperature, the injected fuel can heat quickly and reinforce the advancing burn structure. The fuel is therefore incorporated into the propagating crest as it passes through the region rather than directly modifying the crest at the moment of injection.

This behaviour suggests that the propagating burn in the reduced model functions as a phase-coupled system in which the timing of fuel deposition relative to the crest determines whether the injection acts constructively or destructively. Fuel introduced within the immediate crest interaction region disrupts the local structure of the wave, while fuel introduced outside this region but still within the broader wake of the travelling structure can reinforce the propagation process.

In this interpretation the crest acts as a moving amplification region for thermally activated fusion reactions that entrains fuel deposited within its wake while remaining sensitive to direct perturbations near the peak. The persistence plateau therefore emerges not from the magnitude of the injected fuel alone, but from the phase relationship between injection events and the evolving wave structure.

Although the present model is intentionally simplified, the emergence of this phase-dependent behaviour is qualitatively reminiscent of synchronisation phenomena in nonlinear dynamical systems. The present model does not attempt a formal oscillator description, but the results indicate that phase relationships between fuelling and propagation may represent an important dynamical factor in systems exhibiting travelling burn fronts.

4.2. Structure of the Wake-Lag Window

The lag scans indicate that the persistence plateau occupies a finite angular region relative to the propagating crest rather than extending across the entire domain. The precise extent of this region varies somewhat with the structural collapse criteria used to define crest termination. Collapse occurs when injection events are placed too close to the crest maximum, while sustained propagation emerges once the injection is sufficiently offset from the peak.

Inspection of the temperature field profiles provides physical context for this transition. The crest itself appears as a relatively localised temperature maximum, but the surrounding temperature field extends over a broader angular domain. Injection events occurring within this

region of strong crest influence tend to raise the instantaneous crest temperature while simultaneously shortening the propagation lifetime, indicating that direct perturbation of the crest structure destabilises the wave.

The lag-scan results suggest that the transition to sustained propagation occurs once injection moves beyond this immediate crest interaction region. At these offsets the injected fuel is no longer deposited directly into the crest core but instead enters the broader dynamical envelope of the travelling wave. In this regime the fuel remains phase-coupled to the propagating structure and can be incorporated into the advancing burn front as the crest passes.

Because the system is periodic, this stable fuelling region appears on both sides of the crest in the lag scans. In the context of a travelling wave on a closed toroidal domain, this region can be interpreted as part of the effective wake of the propagating crest, representing the portion of the temperature field that remains dynamically coupled to the advancing structure after the crest maximum has passed. Injection events placed within this wake-coupled region reinforce the wave, whereas injections placed within the immediate crest interaction region destabilise it.

The resulting persistence plateau therefore reflects a balance between two competing effects: the destabilising influence of direct crest perturbation and the reinforcing influence of wake-coupled fuelling. The lag-scan boundary observed in the simulations marks the approximate angular distance at which injection transitions between these two regimes. Sensitivity tests indicate that the precise location of this boundary depends on the structural collapse criterion, although the existence of the wake-aligned persistence region remains robust.

4.3. Model Scope and Limitations

The propagation model used in this study is intentionally reduced in order to isolate the dynamical consequences of phase-dependent fuelling. The system represents a simplified travelling burn structure evolving on a one-dimensional toroidal domain with prescribed heating, transport, and thermalisation processes. Several physical effects that would be present in realistic plasma systems are therefore not included.

In particular, the model does not represent full magnetohydrodynamic plasma behaviour, detailed turbulence-driven transport, multi-species plasma dynamics, or magnetic geometry effects. The representation of fusion heating, thermalisation, and particle exhaust is likewise simplified in order to maintain a tractable propagation model suitable for exploring phase relationships between fuelling events and the evolving burn structure.

As a result, the behaviours identified here should not be interpreted as quantitative predictions for specific reactor configurations. Rather, the model provides a minimal dynamical framework for examining how the timing of localised fuelling interacts with a propagating burn crest.

Within this framework the emergence of a phase-dependent persistence plateau appears to be a robust feature of the propagation dynamics. The behaviour is consistent with phase-dependent propagation mechanisms that also arise in advection-diffusion-reaction systems, suggesting that the effect may reflect a more general coupling between transport, reaction kinetics, and localised fuelling within propagating structures, rather than depending exclusively on fusion-specific physics. Whether similar phase-coupled behaviour persists in more realistic plasma environments will require investigation in higher-dimensional simulations incorporating more complete physical descriptions of transport, magnetic geometry, and plasma stability.

4.4. Implications for Control-Oriented Fuelling

Although the model used in this study is deliberately simplified, the results highlight a potential aspect of fuelling dynamics that is not always emphasized in reduced fusion models: the temporal relationship between fuel injection and evolving burn structures. Fuelling strategies are typically analysed in terms of deposition location, fuelling rate, and total particle input. In systems where a propagating burn structure exists, however, the timing of fuel delivery relative to

that structure may represent an additional control parameter.

The lag-scan results presented here suggest that localised fuelling can interact constructively or destructively with a propagating burn depending on its phase relationship with the crest. Injection events placed too close to the crest tend to perturb the burn structure and shorten propagation lifetime, even when the local temperature temporarily increases. By contrast, injections placed within the broader wake-coupled region of the travelling wave reinforce the propagation and produce the persistence plateau observed in the simulations.

If analogous behaviour were to emerge in more physically complete plasma models, it would imply that fuelling strategies might benefit from synchronisation with the evolving dynamics of the plasma rather than purely static deposition schemes. In systems exhibiting travelling burn structures, injection events timed relative to the propagation of the burn front could potentially influence burn stability in ways not captured by static fuelling models.

One possible framework for exploring this idea is the concept of wake-aligned fuel injection (WAFI), in which fuelling events are timed relative to the propagating structure of the burn wave so that injected fuel enters the wake-coupled region of the crest. The results presented here do not constitute a reactor-scale demonstration of such an approach, but they illustrate how phase-coupled fuelling dynamics can arise even in a minimal propagation model.

Further investigation in higher-dimensional simulations and more physically complete plasma models will be required to determine whether similar synchronisation effects persist beyond the reduced framework considered in this study.

5. Conclusion

This work examined the effect of phase-dependent fuel injection in a reduced propagation model representing a travelling burn structure on a one-dimensional toroidal domain. By scanning the angular offset between a moving temperature crest and localised fuelling events, the simulations reveal a strong dependence of burn persistence on the relative phase between injection and crest position.

Across the wake-lag scans, injections placed close to the crest tend to destabilise the propagating structure, producing rapid collapse despite temporary increases in local temperature. In contrast, injections occurring within a finite wake-coupled region away from the crest lead to sustained propagation and significantly longer burn lifetimes even when the total fuel input is identical. A persistence plateau appears consistently across the explored lag ranges, although its precise boundaries depend on the structural criteria used to define crest collapse.

A series of sensitivity tests and numerical checks indicate that this behaviour is not an artefact of a narrow parameter choice or grid resolution. Variations in termination criteria, transport parameters, and crest speed preserve the overall structure of the wake-lag window, while grid refinement tests demonstrate convergence of the crest profile and associated propagation metrics.

Although the model is intentionally simplified and does not represent full plasma physics, the results highlight a dynamical mechanism in which the timing of fuelling relative to a propagating burn structure can influence propagation stability. If similar behaviour emerges in more complete plasma models, phase-synchronised fuelling strategies such as wake-aligned injection may represent an additional control dimension for systems exhibiting travelling burn dynamics.

While the present results illustrate conceptual dynamics in a one-dimensional reduced model, extrapolation to actual reactor predictions would require fully resolved, higher-dimensional plasma simulations.

Further work will be required to explore whether comparable phase-coupled behaviour persists in higher-dimensional simulations incorporating more realistic transport, magnetic geometry, and stability physics.

6. References

[1] ITER Physics Basis Editors, "ITER physics basis, 1; Overview and summary," *Nuclear Fusion, vol. 39, pages 2137–2174, 1999,* https://doi.org/10.1088/0029-5515/39/12/301.

[2] M. Shimada et al., "Chapter 1: Overview and summary," *Nuclear Fusion, vol. 47, no. 6, pages S1–S17, 2007,* https://doi.org/10.1088/0029-5515/47/6/S01.

[3] A. R. Christopherson, R. Betti, S. Miller, V. Gopalaswamy, O. M. Mannion, D. Cao, "Theory of ignition and burn propagation in inertial fusion implosions," *Physics of Plasmas, vol. 27, 052708, 2020,* https://doi.org/10.1063/1.5143889.

[4] J. D. Lindl, "Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain," *Physics of Plasmas, vol. 2, pages 3933–4024, 1995,* https://doi.org/10.1063/1.871025.

[5] S. Atzeni and J. Meyer-ter-Vehn, The Physics of Inertial Fusion: Beam Plasma Interaction, Hydrodynamics, Hot Dense Matter, International Series of Monographs on Physics, Oxford University Press, 2004, https://doi.org/10.1093/acprof:oso/9780198562641.001.0001.

[6] S. T. O'Neill, B. D. Appelbe, A. J. Crilly, C. A. Walsh, D. J. Strozzi, J. D. Moody, J. P. Chittenden, "Burn propagation in magnetized high-yield inertial fusion," *Physics of Plasmas, vol. 32, 022703, 2025,* https://doi.org/10.1063/5.0242215.

[7] W. W. Heidbrink, "Basic physics of Alfvén instabilities driven by energetic particles in toroidally confined plasmas," *Physics of Plasmas, vol. 15, 055501, 2008,* https://doi.org/10.1063/1.2838239.

[8] C. M. Bishop, J. W. Connor, "Heat-pulse propagation in tokamaks and the role of density perturbations," *Plasma Physics and Controlled Fusion, vol. 32, no. 3, 1990,* https://doi.org/10.1088/0741-3335/32/3/005.

[9] T. S. Hahm, P. H. Diamond, Z. Lin, K. Itoh and S-I Itoh, "Turbulence spreading into the linearly stable zone and transport scaling," *Plasma Physics and Controlled Fusion, vol. 46, pages A323–A333, 2004,* https://doi.org/10.1088/0741-3335/46/5A/036.

[10] S. H. Strogatz, "From Kuramoto to Crawford: Exploring the onset of synchronization in populations of coupled oscillators," *Physica D, vol. 143, no. 1–4, pages 1–20, 2000,* https://doi.org/10.1016/S0167-2789(00)00094-4.

[11] M. C. Cross and P. C. Hohenberg, "Pattern formation outside of equilibrium," *Reviews of Modern Physics, vol. 65, pages 851–1112, 1993,* https://doi.org/10.1103/RevModPhys.65.851.

[12] A. N. Kolmogorov, I. G. Petrovsky, and N. S. Piskunov, "A study of the diffusion equation with increase in the amount of substance and its application to a biological problem," *Bull. Moscow Univ. Math. Mech. 1(6), 1–26 (1937).*

[13] R. B. Bird, W. E. Stewart, and E. N. Lightfoot, Transport Phenomena, 2nd ed., John Wiley & Sons, New York, 2002.

[14] H.-S. Bosch and G. M. Hale, "Improved formulas for fusion cross-sections and thermal reactivities," *Nuclear Fusion, vol. 32, no. 4, pages 611–631, 1992,* https://doi.org/10.1088/0029-5515/32/4/I07.

[15] E. F. Toro, Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction, 3rd ed, Springer, Berlin, 2009, https://doi.org/10.1007/b79761.

[16] P. K. Sweby, “High resolution schemes using flux limiters for hyperbolic conservation laws,” *SIAM Journal on Numerical Analysis, vol. 21, no. 5, pages 995–1011, 1984,* https://doi.org/10.1137/0721062.

Appendix A - Code Availability

The reduced one-dimensional simulation code used in this work is under active development and is not yet prepared for public release.

The numerical implementation used in this study is described in Appendix C.

At present, the code and selected technical materials may be made available for technical review or collaborative research upon reasonable request via: research@chatwoodlabs.com. A curated public release may be considered following documentation, validation, and review of ongoing IP considerations.

The current implementation represents a reduced conceptual model for exploratory study of phase-dependent propagation dynamics and is not intended as a predictive reactor simulation.

Appendix B - Scope and Modelling Context

The modelling framework presented in this work is a reduced exploratory system intended to examine conceptual aspects of phase-coupled fuelling within a propagating burn structure. Several simplifying assumptions are introduced in order to isolate the interaction between localised fuel injection and the dynamics of the travelling crest. These include prescribed crest motion, surrogate fusion heating, simplified transport representations, and idealised injection timing.

These simplifications allow the propagation dynamics to be examined in a controlled setting but should not be interpreted as defining necessary physical mechanisms or operational constraints for practical plasma systems.

The quantitative values and parameter choices reported in this study arise from the specific structure of the reduced model and the numerical environment in which it is implemented. They therefore reflect model-specific conditions rather than universal requirements for systems exhibiting phase-coupled propagation behaviour.

More generally, the phenomena explored here should be understood as one possible realization within a broader class of systems in which localised fuelling interacts with a propagating burn structure. Alternative geometries, transport models, control strategies, or physical mechanisms may produce qualitatively similar synchronisation behaviour within more complete physical descriptions.

The reduced model presented here is therefore intended as a conceptual tool for examining propagation dynamics rather than as a definitive description of any specific reactor configuration.

Appendix C - Numerical Implementation and Diagnostic Definitions

This appendix records the numerical implementation used for the reduced simulations so that the update rules, injection logic, diagnostics, and stopping conditions are explicit. The model is a reduced one-dimensional periodic ring in arbitrary units, intended as a conceptual propagation study rather than a predictive reactor calculation. The evolved fields are thermalised deuterium and tritium densities, fresh non-reacting deuterium and tritium reservoirs, total thermal energy density, and, when exhaust is enabled, helium ash density. Temperature is recovered from the thermalised population as $T = U/(n_D + n_T)$, while fresh fuel contributes to radiative loading immediately but does not participate in fusion until thermalised. Helium ash is tracked separately and is excluded from this reduced temperature closure. Baseline numerical and model settings are summarized in Table 1; the present appendix records the implemented update rules, diagnostics, and stopping logic in explicit form.

At each timestep, the thermalised DT reaction rate is evaluated as:

$$R_{\mathrm{fus}} = n_D n_T \langle \sigma v \rangle_{DT}(T),$$

where $\langle \sigma v \rangle_{DT}$ is taken from the Bosch-Hale fit used in the code [14]. Fusion power is then represented in model units as $P_{\mathrm{fus}} = G R_{\mathrm{fus}}$, with G the dimensionless fusion-gain scaling coefficient. A fraction α of this power is redistributed as alpha heating through a normalized forward-biased kernel, while radiation losses are represented by:

$$P_{\mathrm{rad}} = C_{\mathrm{rad}}\, n_e^2 \sqrt{T},$$

with $n_e = n_D + n_T + n_D^{\mathrm{fresh}} + n_T^{\mathrm{fresh}}$. This means fresh injected fuel increases radiative loading before it becomes burn eligible.

Transport is implemented on the periodic ring with resolution-consistent discretization. Diffusion uses the discrete Laplacian $[1, -2, 1]/\Delta\phi^2$, so that the physical diffusion strength is preserved under grid refinement. Advection is computed with a MUSCL-Hancock scheme using a minmod TVD limiter [15] [16]. Energy density and fresh-fuel reservoirs are advected at the nominal crest speed, while the thermalised deuterium and tritium populations are advected at $v_{\mathrm{therm}} = \epsilon_{\mathrm{therm}} v_{\mathrm{crest}}$. Alpha-heating kernel width and forward bias are likewise scaled with resolution so that their angular extent remains fixed in physical angle rather than in grid cells. The timestep is scaled with resolution from a CFL-safe base value, and the code enforces advection and diffusion stability checks before the run begins.

Fresh fuel is represented by separate non-reacting reservoirs n_D^{fresh} and n_T^{fresh}. Thermalisation is modelled as a first-order transfer:

$$\frac{dn_D^{\mathrm{fresh}}}{dt} = -\lambda_{\mathrm{therm}} n_D^{\mathrm{fresh}}, \frac{dn_T^{\mathrm{fresh}}}{dt} = -\lambda_{\mathrm{therm}} n_T^{\mathrm{fresh}},$$

with the transferred mass added to the thermalised n_D and n_T populations. No additional energy is added during thermalization. Instead, the enthalpy of the injected fuel is added at injection time, and the subsequent increase in thermalised particle number lowers or preserves the mixed temperature through the relation $T = U/n$. This delayed-fusion-eligibility model is the mechanism by which fuel can be deposited at one phase of the orbit and become burn-eligible only after advection and equilibration.

Fuel injection is handled by a set of fixed injectors distributed uniformly around the ring. The injector count is scaled with resolution to preserve injector angular density, and the per-injector fuel amount is scaled

inversely so that the total fuel delivered per cycle remains constant across resolutions. Each injector may fire at most once per orbit. In wake mode, the target angle is defined by:

$$\phi_{\text{target}} = \phi_{\text{crest}} + \Delta\phi_{\text{lag}},$$

with positive lag corresponding to injection behind the crest and negative lag to injection ahead of it. A wake-mode firing is accepted only when the injector lies within an angular acceptance window of $\max(1°, \Delta\phi_{\text{cell}})$from the target angle. In asynchronous mode, there is no angular alignment condition; injectors fire from their own fixed locations once the once-per-orbit gate is satisfied. Injected fuel is deposited as a normalized Gaussian plume in angle, centered on the target angle in wake mode and on the injector angle in asynchronous mode, so that total injected mass and injected enthalpy are conserved by construction.

The code supports both fixed-temperature and local-temperature injection. In fixed mode, each event uses the prescribed injection temperature T_{fuel}. In local mode, the injected enthalpy is based on the Gaussian-weighted local plasma temperature plus an optional modifier, subject to clipping within the simulation bounds. The baseline parameter set reported in the manuscript uses fixed-temperature injection, while the code also retains the local-temperature option for auxiliary-heating-style tests. In all cases, injected mass is first placed in the fresh, non-reacting reservoir and only later transferred to the reacting population by the thermalisation rule described above.

When exhaust is enabled, DT reactions produce helium ash and consume one deuterium and one tritium from the thermalised reacting pool. Excess deuterium and tritium above their target densities are then removed on the characteristic pumping timescale $\tau_{\text{pump,DT}}$, while helium ash is removed on $\tau_{\text{pump,He}}$. Exhaust can be applied either globally or through a normalized Gaussian divertor mask. Energy is removed consistently with the exhausted particles through $dU = -T\, dn_{\text{pumped}}$, after which n, T, and U are resynchronised. The code also tracks cumulative alpha input, radiation losses, injected fuel enthalpy, and exhaust losses to monitor energy closure throughout the run.

The propagating crest is tracked by the location of the maximum temperature, $\phi_{\text{crest}} = \arg\max T$, with periodic unwrapping used to construct a continuous orbit count. This orbit count is used both for once-per-orbit injection gating and for the forcing-clock diagnostics. For the principal CSV trace, the reported crest amplitude is the mean temperature within a fixed physical window of $\pm 3°$around the hottest cell. Temperature contrast is defined as crest amplitude minus the mean ring temperature. Instantaneous crest velocity is obtained from successive crest positions and written to the diagnostic CSV together with crest position, temperature contrast, fusion power, radiative loss rate, injected fuel total, and closure errors. Each run also writes a metadata header containing the numerical parameters used.

In addition to the raw timestep diagnostics, the analysis uses fuel-normalised crest-sustainment measures derived from the crest-local temperature diagnostic. Let $A_{\text{crest}}(t)$denote the crest amplitude, defined as the mean temperature within a fixed $\pm 3°$window about the hottest cell. The integrated efficiency reported in Table 6 is:

$$\eta_{\text{int}} = M_{\text{fuel}}^{-1} \int_{t_{\text{start}}}^{t_{\text{end}}} A_{\text{crest}}(t)\, dt,$$

where $M_{\text{fuel}}$is the cumulative injected fuel mass over the run. The burn-time efficiency is:

$$\eta_{\mathrm{burn}} = M_{\mathrm{fuel}}^{-1} \int_{t_{\mathrm{start}}}^{t_{\mathrm{end}}} \mathbf{1}\left[T_{\mathrm{peak,crest}}(t) > 10\ \mathrm{keV}\right] dt,$$

where $T_{\mathrm{peak,crest}}$is the crest-local peak temperature. The closure error reported in Table 6 is the maximum relative energy-closure residual attained during the run. For Figure 2, the quantity labelled crest sustainment per unit fuel injected refers to the instantaneous fuel-normalised crest amplitude, $A_{\mathrm{crest}}/M_{\mathrm{fuel}}$, evaluated from the timestep diagnostics.

A simulation is terminated when any of three collapse criteria is met: (i) the measured crest speed falls below the prescribed stall threshold, (ii) the peak temperature falls below the temperature threshold, or (iii) the temperature field becomes too flat, quantified by a peak-to-mean temperature ratio below the flatness threshold. In the reported baseline configuration these thresholds are 0.5 rad per model time unit, 10.5 keV, and 1.02 respectively. The collapse test is applied during the run using periodic speed measurements and is repeated once at the end of the run to avoid missing a terminal event in the final measurement interval.

Figures 4-7 were generated from the timestep CSV and per-event injection log written by the simulation. Inter-fire cadence is computed for each injector as the difference in successive firing orbit counts. The forcing diagnostic stores both the actual crest orbit count and the scheduler's target orbit count at each fire event. The residual plotted in Figure 7 is therefore:

$$\Delta N_{\mathrm{res}} = N_{\mathrm{orbit}} - N_{\mathrm{target}},$$

evaluated only at injection events. A flat but nonzero residual indicates a stable timing offset between crest motion and forcing, whereas a systematic slope in the residual would indicate drift of the forcing clock relative to the crest. Rolling mean and ±σ envelopes in Figures 4, 5, and 7 are applied in post-processing with sliding windows proportional to the duration of the analysed interval, and the Figure 6 histogram is formed from the pooled inter-fire intervals across injectors.

The present implementation is a deliberately reduced 1D mechanism model in arbitrary units, designed to isolate phase-dependent fuelling dynamics; source strengths, thermalisation, and collapse criteria are therefore model-defining closures rather than first-principles reactor predictions.