





A Fast and Accurate Method to Capture the Solar Corona/Transition Region Enthalpy Exchange

C. D. Johnston¹  and S. J. Bradshaw² 

¹School of Mathematics and Statistics, University of St Andrews, St Andrews, Fife, KY16 9SS, UK; cdj3@st-andrews.ac.uk

²Department of Physics and Astronomy, Rice University, Houston, TX 77005, USA; stephen.bradshaw@rice.edu

Received 2019 February 3; revised 2019 March 3; accepted 2019 March 3; published 2019 March 15

Abstract

The brightness of the emission from coronal loops in the solar atmosphere is strongly dependent on the temperature and density of the confined plasma. After a release of energy, these loops undergo a heating and upflow phase, followed by a cooling and downflow cycle. Throughout, there are significant variations in the properties of the coronal plasma. In particular, the increased coronal temperature leads to an excess downward heat flux that the transition region (TR) is unable to radiate. This generates an enthalpy flux from the TR to the corona, increasing the coronal density. The enthalpy exchange is highly sensitive to the TR resolution in numerical simulations. With a numerically underresolved TR, major errors occur in simulating the coronal density evolution and, thus, the predicted loop emission. This Letter presents a new method that addresses the difficulty of obtaining the correct interaction between the corona and corona/chromosphere interface. In the TR, an adaptive thermal conduction approach is used that broadens any unresolved parts of the atmosphere. We show that this approach, referred to as TRAC, successfully removes the influence of numerical resolution on the coronal density response to heating while maintaining high levels of agreement with fully resolved models. When employed with coarse spatial resolutions, typically achieved in multidimensional MHD codes, the peak density errors are less than 3% and the computation time is three orders of magnitude faster than fully resolved field-aligned models. The advantages of using TRAC in field-aligned hydrodynamic and multidimensional magnetohydrodynamic simulations are discussed.

Key words: hydrodynamics – magnetohydrodynamics (MHD) – Sun: chromosphere – Sun: corona – Sun: flares – Sun: transition region

Supporting material: animation

1. Introduction

One of the main difficulties encountered when using field-aligned hydrodynamic models to study the physics of magnetically closed coronal loops is the need to implement a grid that fully resolves the steep and dynamic transition region (TR) (see, e.g., Bradshaw & Cargill 2013, hereafter BC13). In a static coronal loop, the energy balance in the TR is between a downward heat flux from the corona and optically thin radiative losses to space. Defining the temperature length scale as $L_T = T/|dT/ds|$, where s is a coordinate along the magnetic field, such static loops can have L_T of order 1 km in the TR. During impulsive heating, the heat flux from the corona can be enhanced by orders of magnitude, so that L_T becomes very small, of order 100 m and less in some cases.

Properly resolving these small length scales is essential in order to obtain the correct coronal density in response to heating (BC13); otherwise, the downward heat flux “jumps” over an underresolved TR to the chromosphere, where the incoming energy is then strongly radiated. BC13 showed that major errors in simulating the coronal density evolution were likely with a lack of spatial resolution.

However, because the thermal conduction timescale across a grid cell scales as $(\Delta s)^2$ and $T^{-5/2}$ (where Δs is the grid cell width), cells in the cooler lower TR must become dramatically smaller to resolve the temperature gradient that supports the incoming heat flux. In addition, numerical stability requires that the minimum timescale across the entire grid is temporally resolved. Thus, locally satisfying $\Delta s < L_T$ in the TR implies long computation times for fully resolved field-aligned simulations. The problem is more severe in

three-dimensional (3D) magnetohydrodynamic (MHD) models where computational resources place significant constraints on the achievable resolutions. Therefore, other approaches need to be found.

One such approach has been developed by Linker et al. (2001), Lionello et al. (2009), and Mikić et al. (2013) for use in global 3D MHD simulations. This technique modifies the temperature dependence of the parallel thermal conductivity $\kappa_{\parallel}(T)$ and radiative emissivity $\Lambda(T)$ below a fixed temperature value (T_c) in the TR. It has been demonstrated that such modifications artificially broaden the TR for $T < T_c$ without significantly changing the coronal properties of the loop. While this is a particularly attractive property the method still requires sufficiently high resolution to properly resolve the small length scales in the TR above T_c . Furthermore, the characteristic temperature of the TR can change dramatically during the heating and cooling cycles, making it unclear how suitable a fixed value of T_c is for capturing the dynamic evolution of coronal loops that are impulsively heated.

Another approach has been developed by Johnston et al. (2017a, 2017b), who treated the (unresolved) TR as a discontinuity across which energy is conserved through the imposition of a jump condition. No attempt is made to properly resolve the TR. Instead, the method is employed with coarse spatial resolutions that are supported by dynamically locating the top of the unresolved transition region (UTR). This approach gives good agreement with fully resolved HYDRAD simulations (Bradshaw & Mason 2003; Bradshaw & Cargill 2006, BC13) while significantly speeding up the computation time (e.g., Johnston et al. 2019).

This Letter proposes a new approach for modeling the TR that incorporates the strengths of both approaches. Principally, the ability to dynamically identify the UTR to prescribe an adaptive T_c and then broaden the length scales in this unresolved region only, in order to eliminate the need for highly resolved numerical grids. The outcome is an extremely powerful method that (1) removes the influence of numerical resolution on the coronal response to heating and (2) accurately predicts the results of properly resolved field-aligned models (e.g., BC13) when employed with the coarser spatial resolutions achieved by multidimensional MHD codes.

2. Numerical Model and Experiments

2.1. Numerical Model

We solve the field-aligned hydrodynamic equations using the HYDRAD code (Bradshaw & Mason 2003; Bradshaw & Cargill 2006, BC13) run in single fluid mode. HYDRAD uses adaptive regridding to ensure adequate spatial resolution, with the grid being refined such that cell-to-cell changes in the temperature and density are kept between user defined values (taken as 5% and 10% here), where possible. The largest grid cell in all of our calculations has a width of 10^6 m (1000 km) and each successive refinement splits the cell into two. Thus, a refinement level of RL leads to cell widths decreased by $1/2^{\text{RL}}$. The maximum value of RL is limited to 14, corresponding to a grid cell width of 61 m in the most highly resolved parts of the TR.

2.2. TR Adaptive Conduction

While it is possible to employ such high resolution grids to properly resolve the TR in field-aligned codes, this approach is unlikely to be a viable way of running multidimensional MHD simulations. Therefore, we have developed a new method that addresses this difficulty of obtaining adequate spatial resolution in numerical simulations by modeling the transition region using adaptive conduction (TRAC) coefficients that act to broaden any unresolved parts of the atmosphere. This treatment of thermal conduction in the TR, referred to as the TRAC method, relies on a dynamic capability to track resolved and modify unresolved conductive fluxes.

2.2.1. Identification of an Adaptive Cutoff Temperature

The implementation of TRAC is comprised of two main parts. The first part is to identify the maximum temperature of any unresolved grid cells in the numerical domain and use the calculated value to prescribe an adaptive cutoff temperature (T_c). This is done by using an algorithm that is based on the method employed by Johnston et al. (2017a, 2017b) for locating the top of the UTR in the jump condition method as follows.

We define the temperature length scale as,

$$L_T(T(s)) = \frac{T}{|dT/ds|}, \quad (1)$$

and the resolution in a simulation is given by the local grid cell width,

$$L_R(s) = \Delta s, \quad (2)$$

where s is the spatial coordinate along the magnetic field.

Using these definitions, the cutoff temperature is defined as the maximum temperature that violates the resolution criteria of Johnston et al. (2017a, 2017b),

$$T_c = \max(T(s)) \left| \frac{L_R(s)}{L_T(s)} > \delta = \frac{1}{2}, \quad (3)$$

which corresponds to not having multiple grid cells across the temperature length scale (i.e., unresolved temperature gradients).

An upper bound for the cutoff temperature is set as 20% of the peak coronal temperature in the loop at the time when T_c is evaluated (though the results are only weakly dependent on the maximum temperature fraction),

$$T_{\text{max}} = 0.2 T_{\text{peak}}, \quad (4)$$

and a lower bound set as the temperature value of the isothermal chromosphere,

$$T_{\text{min}} = T_{\text{chrom}}. \quad (5)$$

In this Letter the lower bound is taken as $T_{\text{chrom}} = 2 \times 10^4$ K. Employing these definitions, we dynamically adjust T_c with the criteria that it should satisfy,

$$T_{\text{min}} \leq T_c \leq T_{\text{max}}. \quad (6)$$

Hence, the cutoff temperature is adaptive in identifying the UTR, with the value of T_c used in the method changing in response to coronal heating and cooling.

2.2.2. Broadening the UTR

The second part of TRAC is to broaden the steep temperature and density gradients in the UTR. This is achieved using the approach developed by Linker et al. (2001), Lionello et al. (2009), and Mikić et al. (2013).

Below the adaptive cutoff temperature (T_c), the parallel thermal conductivity (κ_{\parallel}) is set to a constant value,

$$\kappa_{\parallel}(T) = \kappa_0 T^{5/2}, \quad \forall T \geq T_c; \quad (7)$$

$$\kappa_{\parallel}(T) = \kappa_{\parallel}(T_c), \quad \forall T < T_c, \quad (8)$$

and the radiative loss rate (Λ) is modified to preserve $\Lambda(T) \kappa_{\parallel}(T)$,

$$\Lambda(T) = \Lambda(T), \quad \forall T \geq T_c; \quad (9)$$

$$\Lambda(T) = \Lambda(T) \left(\frac{T}{T_c} \right)^{5/2}, \quad \forall T < T_c. \quad (10)$$

Increasing the parallel thermal conductivity and decreasing the radiative loss rate, at temperatures below T_c , has the desired effect of broadening the temperature length scales in the UTR. This helps TRAC prevent the heat flux jumping across the unresolved region while maintaining accuracy in the properly resolved parts of the atmosphere. (The broadening effect will be explained in greater detail in forthcoming work.) Furthermore, we note that the formulation of TRAC (1) makes no assumptions about the spatial resolution in a simulation and (2) the parallel thermal conductivity reduces to the classical Spitzer-Harm value when the TR is properly resolved.

2.3. Validation Experiments

The effectiveness of the TRAC method to obtain the correct interaction between the corona and TR, in response to rapid heating events, is investigated by considering two impulsive

coronal heating events, comprising short and long pulses that last for a total duration of 60 s and 600 s, respectively. The temporal profile of the heating pulses is triangular with a peak value of $Q_H = 2 \times 10^{-2} \text{ J m}^{-3} \text{ s}^{-1}$, while the spatial profile is uniform along the loop. We release the energy in a coronal loop of total length 100 Mm, which includes a 10 Mm chromosphere attached at the base of each TR. Thus, the total energy injected into the coronal part of the loop is $4.8 \times 10^7 \text{ J m}^{-2}$ ($4.8 \times 10^8 \text{ J m}^{-2}$) in the 60 s (600 s) heating pulse simulations. These heating conditions are representative of reasonably powerful flares and present a challenge to resolving the TR.

For each simulation, the main assessment of the performance of TRAC is a comparison with the results from two alternative methods that are commonly used to treat thermal conduction. Each method is applied with spatial resolutions that cover several orders of magnitude, ranging from those required for 3D MHD codes to run in a realistic time to those that fully resolve the TR but are achievable only in field-aligned models. The first of these methods is the classical Spitzer-Harm heat flux formulation, while the second is the approach developed by Lionello et al. (2009), which artificially broadens the TR below a fixed specified temperature. For brevity and clarity we define these two conduction methods as SH and L09, respectively. We note that the L09 method uses the broadening technique outlined in Section 2.2.2 but the modifications are applied below a fixed cutoff temperature taken as $T_c = 250,000 \text{ K}$ (which is a typical value used by Lionello et al. 2009 and Mikić et al. 2013).

Using the same set of parameter values but employing the three different conduction methods (SH, L09, and TRAC), we repeated each simulation for $RL = [0, 1, 3, 5, 7, 9, 11, 13, 14]$ refinement levels to create a group of simulations run for each conduction method.

All of the simulations start from the same initial conditions. These are calculated using the SH parallel thermal conductivity and unmodified radiative loss rate. At $t = 0$ (the initial conditions) a small spatially uniform background heating term (Q_{bg}) is present. This gives a starting temperature of order 1 MK. However, we note that Q_{bg} is switched off thereafter and the total energy in the initial conditions is negligible compared with the energy released into the loop during the heating pulses.

3. Results

3.1. Coronal Response to Heating

Figure 1 shows the temporal evolution of the coronal averaged temperature (T) and density (n) and the corresponding T versus n phase space plot (shown on a log–log scale) for the 60 s heating pulse simulation. The coronal averages are calculated by spatially averaging over the uppermost 50% of the loop. Each conduction method is identified with a specific color and their results are shown in separate rows. The red curves correspond to the SH method (first row), the purple curves the L09 method (second row) and the blue curves represent the TRAC method (third row).

In the panels of each method, each curve corresponds to a simulation run with a different value of maximum refinement level. Taking the SH simulations as an example, the line styles associated with the different refinement levels are shown in the figure legend on the temperature plot, and indicate that RL increases as one moves upwards from the lowest curve

($RL = 0$) to the highest curve ($RL = 14$) in the density plot. These simulations are identical in all respects except for the value of RL . Note that the results for $RL = 11$ and 13 are not shown in the SH panels. In addition, $RL = 11, 13,$ and 14 are not shown in the L09 and TRAC panels, where instead the SH simulation run with $RL = 14$ (solid red curves) is shown. This SH simulation ($RL = 14$) is properly resolved and used as a benchmark solution.

The coronal response of the benchmark simulation (SH with $RL = 14$), to the 60 s heating pulse, follows the standard evolution of an impulsively heated loop as described in the literature (Bradshaw & Cargill 2006; Cargill et al. 2012a, 2012b, 2015; Bradshaw & Klimchuk 2015; Reale 2016). The rise in temperature is followed by the density increasing due to the ablation process (Antiochos & Sturrock 1978; Klimchuk et al. 2008), then, after the time of the density peak, the loop cools by radiation and drains by a downward enthalpy flux to the TR (Bradshaw & Cargill 2010a, 2010b). The temperature cools below the starting value of 1 MK at around 2100 s, and then continues to decrease toward the chromospheric value of $2 \times 10^4 \text{ K}$ while the density is evacuated from the loop.

Consistent with BC13, the coronal density evolution in the group of SH simulations is strongly dependent on the spatial resolution, requiring grid cell widths of at least 488 m ($RL \geq 11$) for convergence. These resolution requirements are slightly weaker in the L09 simulations, with convergence seen for minimum grid cell widths of 7.81 km or finer ($RL \geq 7$). However, the L09 method achieves this relatively modest relaxation of the resolution dependence at the expense of accurately modeling the loop evolution below the specified cutoff temperature ($T_c = 250,000 \text{ K}$).

It is therefore striking that the TRAC solutions are only weakly dependent on the spatial resolution, while maintaining high levels of accuracy throughout the full coronal response range. Grid cell widths of 125 km ($RL = 3$, medium dashed curves) are sufficient to observe convergence to the properly resolved SH solution that employed TR grid cell widths of 61 m. The TRAC solution computed with $RL = 3$ correctly captures the interaction between the corona and chromosphere during the ablation phase, up to the time of peak density where the error is less than 3% and the subsequent decay phase. Even the density oscillations (that are characteristic of the short heating pulse imposed, e.g., Reale 2016) and the global cooling of the loop down to $2 \times 10^4 \text{ K}$, are both correctly captured. This demonstrates that adaptively broadening unresolved parts of the TR removes the influence of numerical resolution on the coronal density response to impulsive heating.

Figure 2 shows the results for the 600 s heating pulse simulations. Releasing the energy for an extended period of time but with the same peak heating rate has two main consequences. First, the density oscillations associated with the sloshing of the coronal plasma disappear. Second, the differences between the density curves that represent the various SH and L09 simulations become more severe.

Considering the refinement level of $RL = 3$ as an example, these conduction methods have peak density errors of 70% (SH) and 20% (L09) when benchmarked against the properly resolved simulation (SH with $RL = 14$), increased from discrepancies of 55% (SH) and 15% (L09) for the shorter heating pulse. Consequently, the convergence requirement on the spatial resolution becomes stricter for the L09 method and remains as

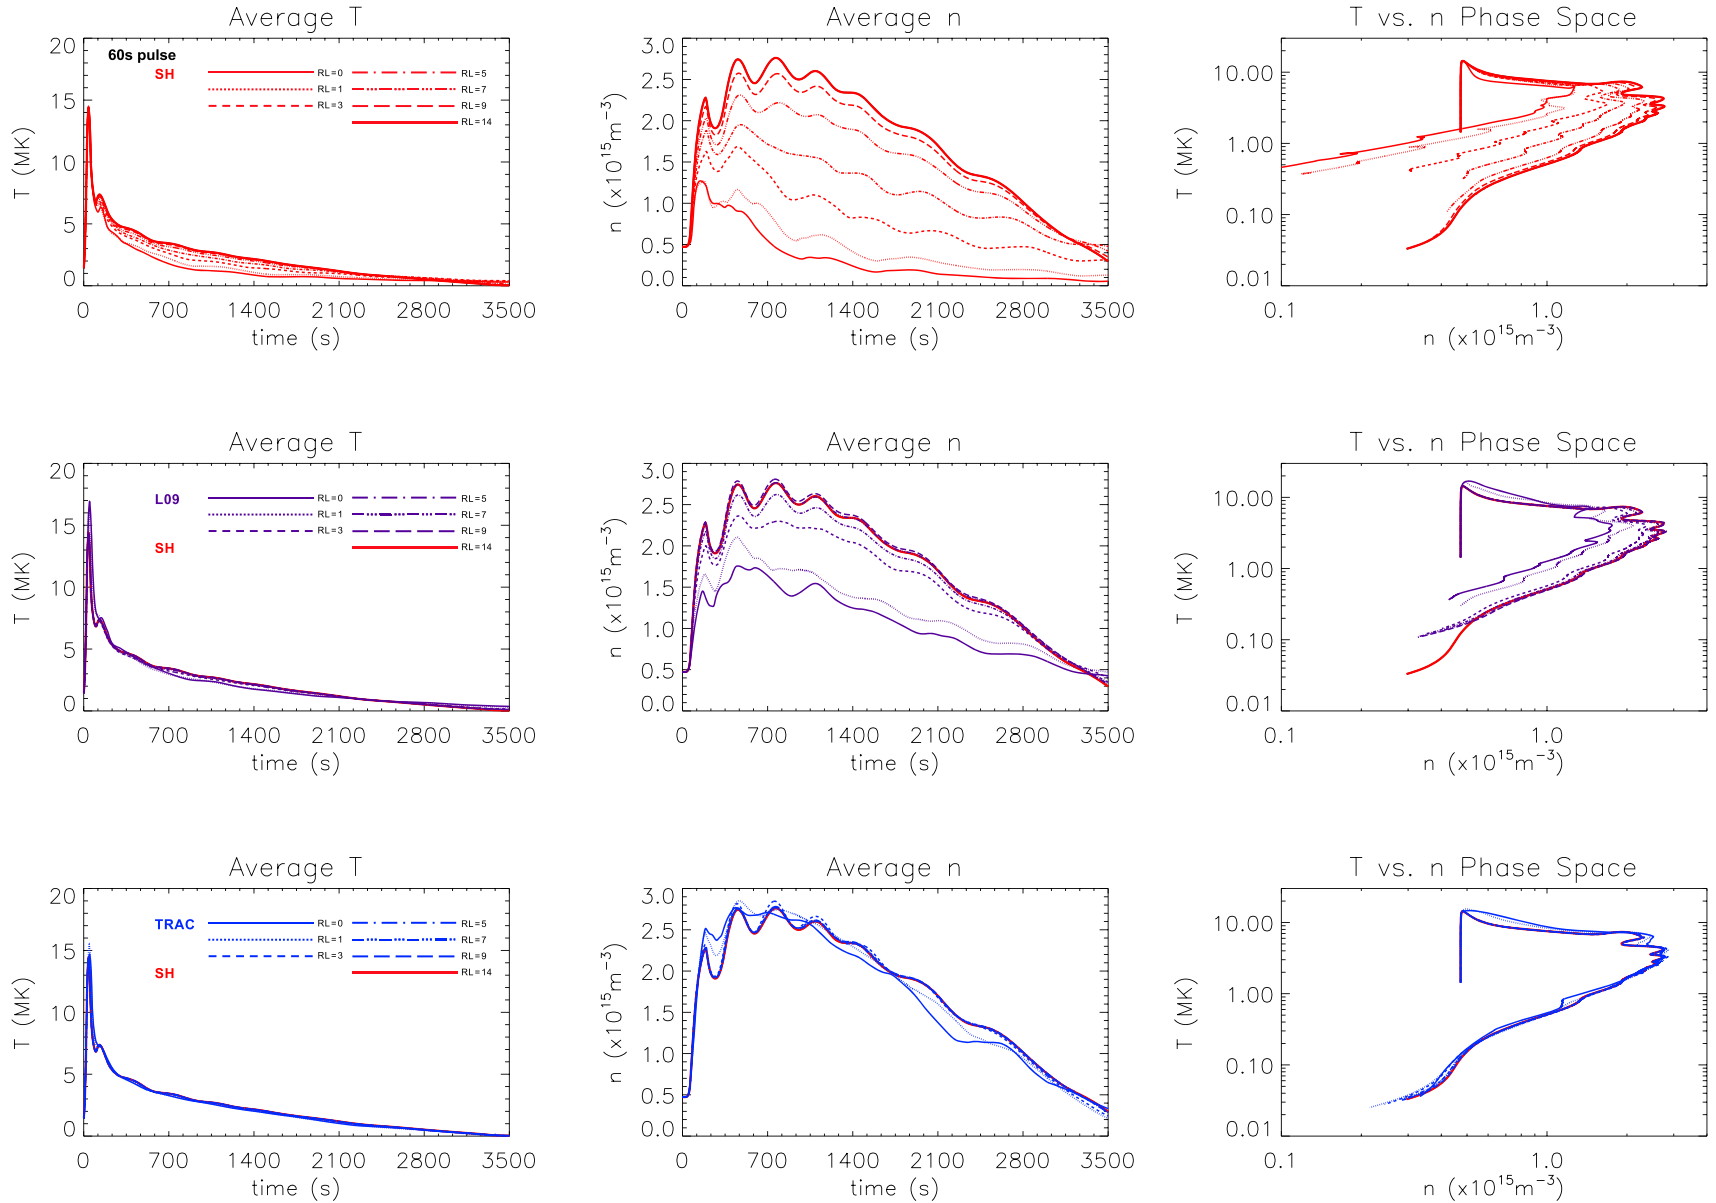


Figure 1. Results for the 60 s heating pulse simulations. The panels show the coronal averaged temperature (left-hand column) and density (central column) as functions of time, and the temperature vs. density phase space plot (right-hand column). The various curves represent different values of RL , which converge as RL increases (higher spatial resolution is associated with larger RL). Rows 1–3 correspond to simulations run with the Spitzer-Harm (SH), Lionello et al. (2009), and TRAC conduction methods, respectively. The lines are color-coded in a way that reflects the conduction method used. Note that the properly resolved SH solution ($RL = 14$) is also shown in the L09 and TRAC panels.

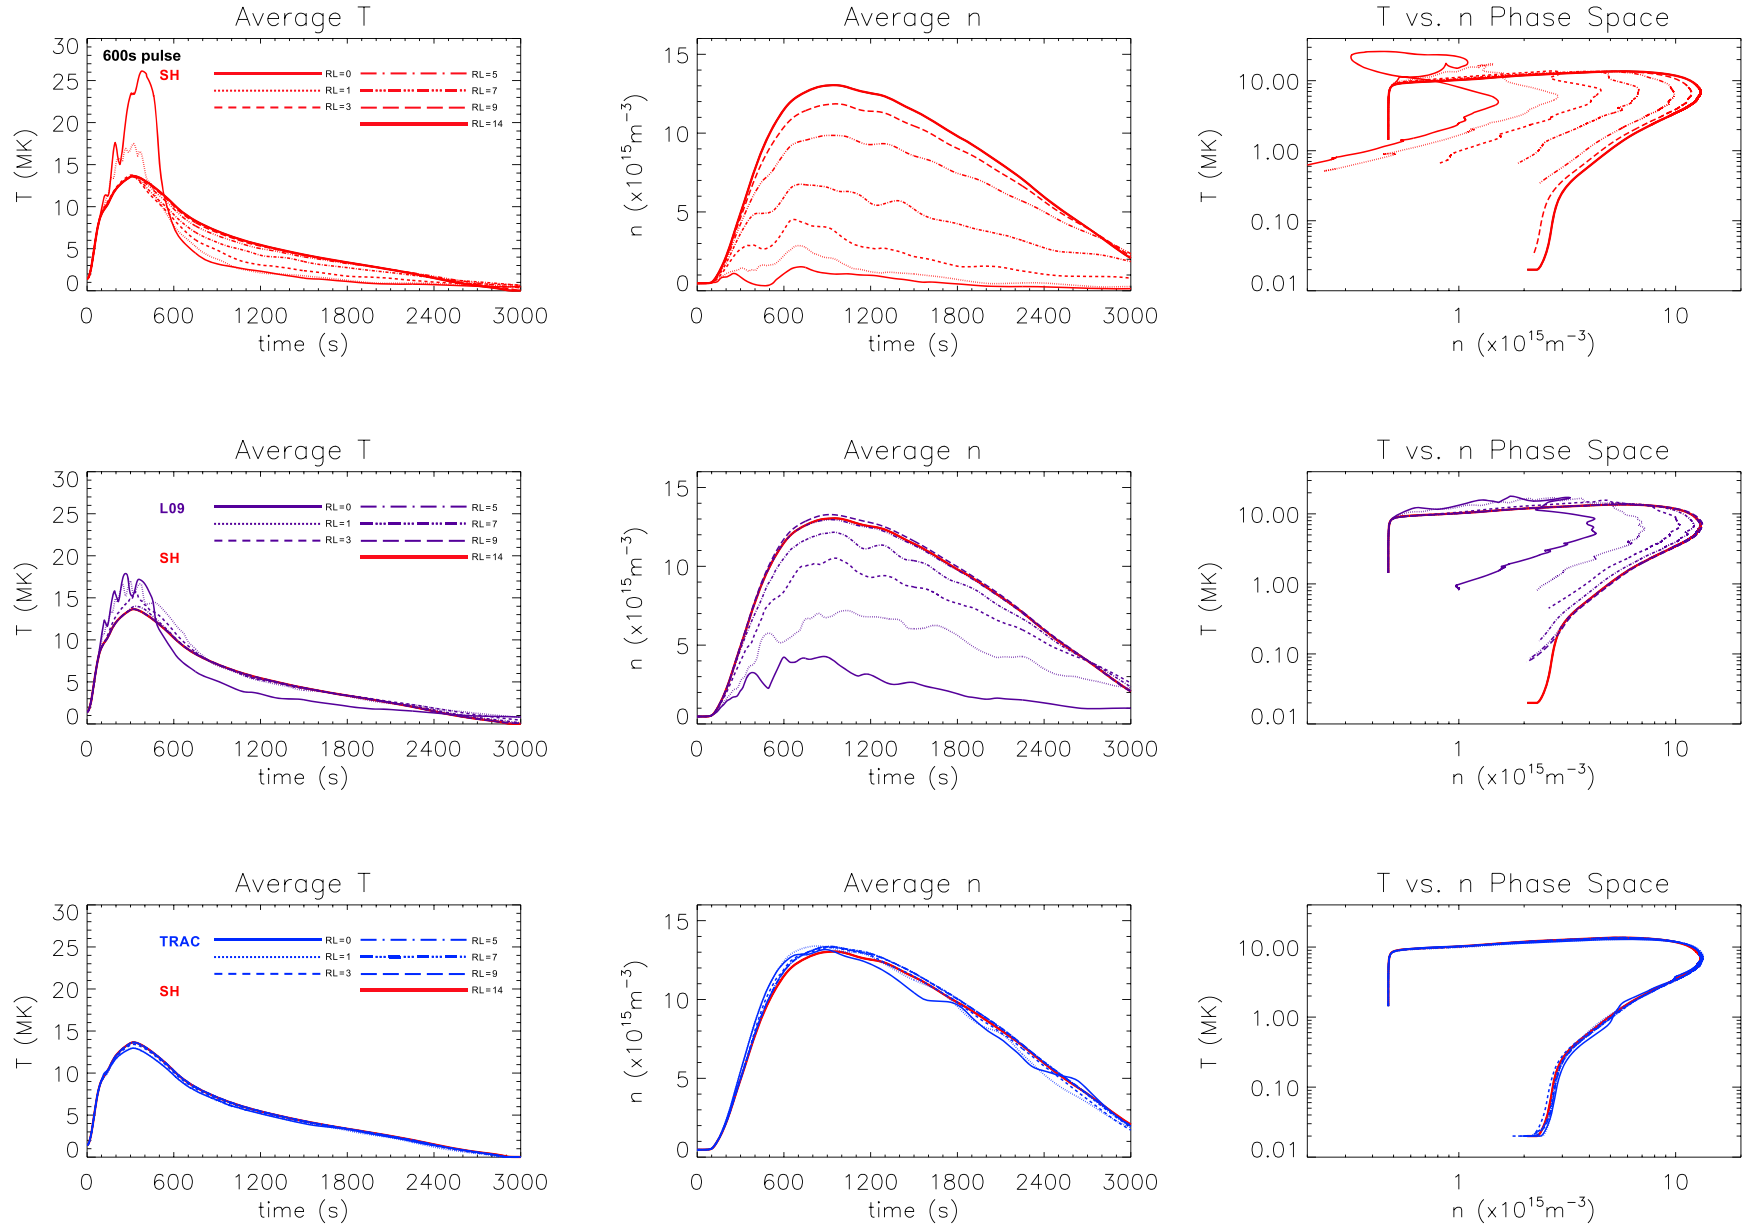


Figure 2. Results for the 600 s heating pulse simulations. Notation is the same as that in Figure 1.

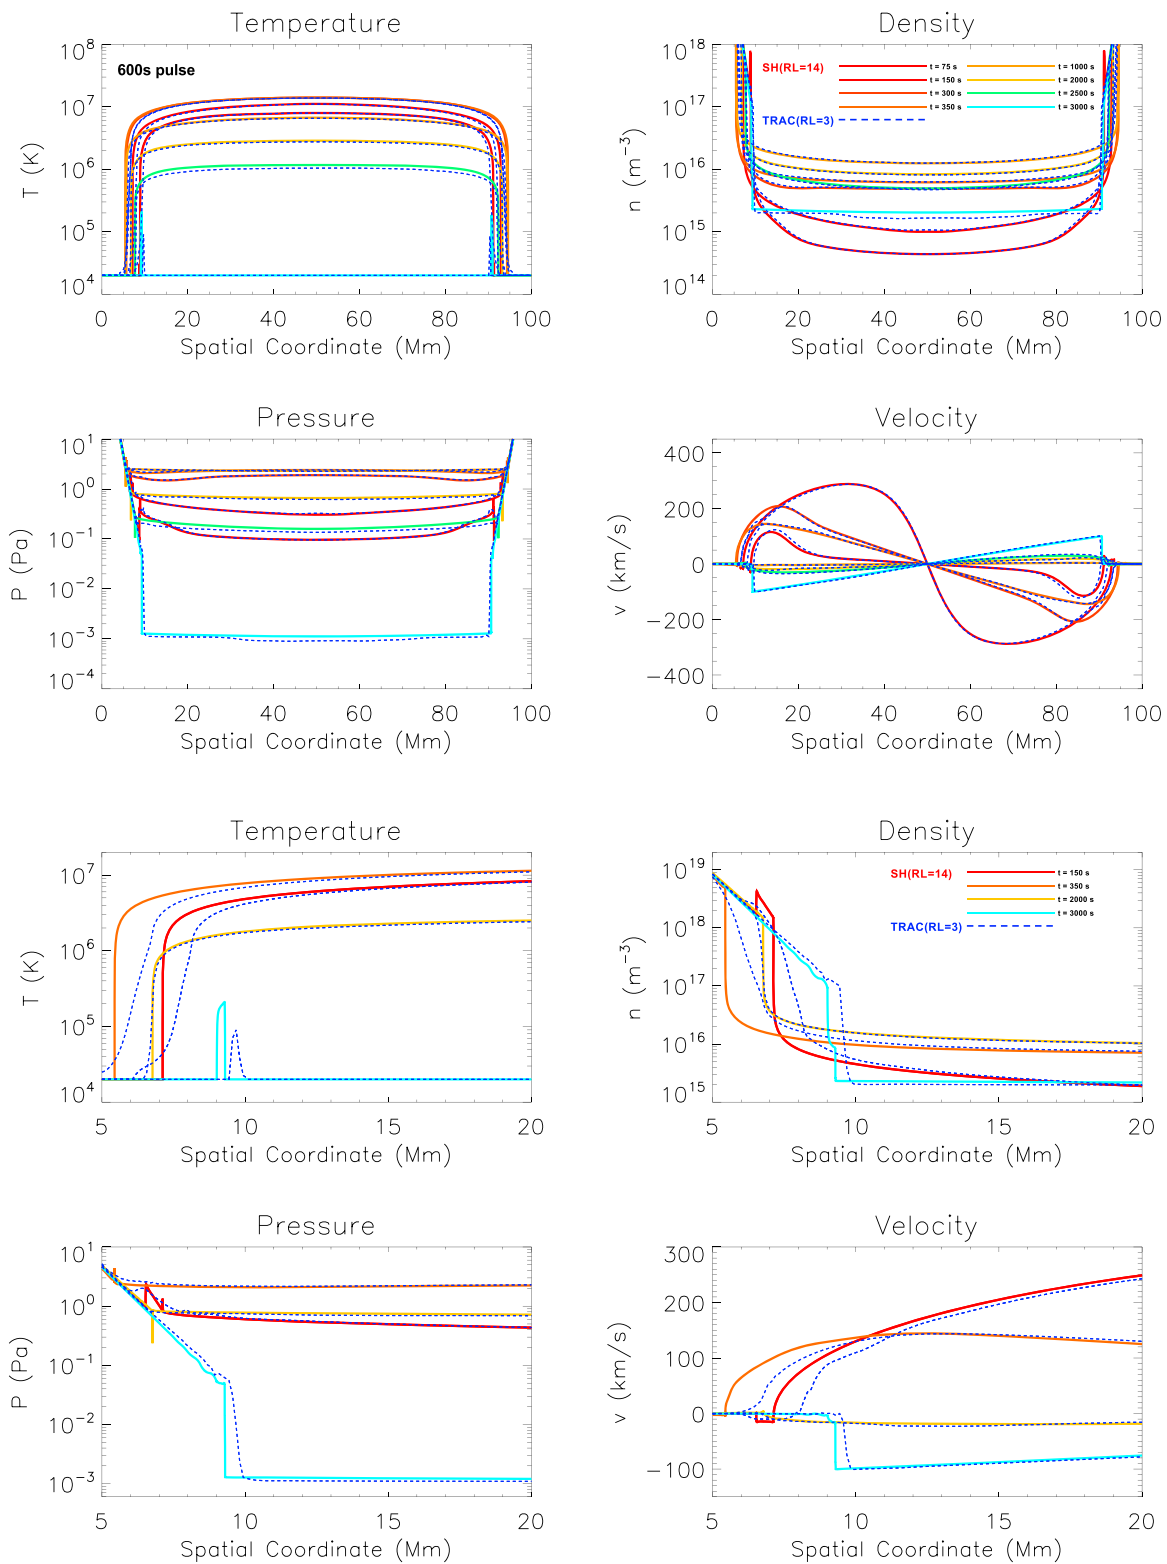


Figure 3. Results for the 600 s heating pulse simulations. Starting from the top left, the upper four panels show time ordered snapshots of the temperature, density, pressure, and velocity as functions of position along the loop for times during the ablation and decay phases. The lower four panels show an enlargement about the transition region of a subset of these snapshots. The various solid curves represent the properly resolved SH solution ($RL = 14$) at different times (with the lines color-coded in a way that reflects the temporal evolution) and the dashed blue curves correspond to the TRAC solution that is computed with $RL = 3$ at these times. Animations of the 60 and 600 s heating pulse simulations can be viewed in the online version of this article. These compare the full time evolution of the T , n , p , and v profiles of the SH simulations run with $RL = 14$ and the TRAC simulations run with $RL = 3$.

(An animation of this figure is available.)

severe for the SH method. Grid cell widths of 488 m ($RL \geq 11$) and 1.95 km ($RL \geq 9$) were necessary for the solutions to converge in the SH and L09 simulations, respectively.

On the other hand, the TRAC solutions once again converge at $RL = 3$ (125 km grid cells) with a peak density error of less than 2%, indicating that this treatment of the TR gives rise to a coronal response that can accurately follow the complete coronal heating and cooling cycle with coarse spatial resolutions, irrespective of the heating duration and resulting density regime.

In particular, TRAC eliminates the need for the very short time steps that are imposed by a highly resolved numerical grid. This enables the method to accurately capture the corona/TR enthalpy exchange significantly faster than fully resolved models. A typical speed-up is between two and three orders of magnitude. For example, the computation time of the properly resolved SH simulation ($RL = 14$) is 3.5 days for the 600 s heating pulse, while the run time for the TRAC simulation computed with $RL = 3$ is only 6 minutes.

3.2. Global Evolution

The main question that needs to be addressed in understanding these results is why are the TRAC simulations so successful in describing the coronal response to heating with such large grid cell widths? A comprehensive answer requires a detailed assessment of how the TRAC modifications to the parallel thermal conductivity and radiative loss rate affect the local energy balance and subsequent dynamics in the TR, coupled together with the resulting global evolution of the loop. Here we concentrate just on the latter. The former will be presented in future work.

Figure 3 shows the time evolution of the global temperature, density, pressure, and velocity profiles for the 600 s heating pulse simulations. Each solid curve represents a different snapshot from the properly resolved SH simulation ($RL = 14$) and the dashed blue curves imposed on top are the corresponding snapshots from the TRAC simulation computed with 125 km grid cells ($RL = 3$). (The comparison with an underresolved SH solution (e.g., SH run with $RL = 3$) has been previously discussed in BC13 and Johnston et al. (2017a) and will not be considered further.)

First we focus on the evolution of the pressure. Of particular importance is the formation of the pressure gradients in the TR. Adaptively broadening the temperature and density profiles in the unresolved parts of the TR prevents the downward heat flux from jumping across the unresolved region. This ensures that the incoming energy goes into increasing the gas pressure locally, rather than being lost due to artificially high radiative losses (e.g., BC13; Johnston et al. 2017a, 2017b). The scale of the broadening is relatively small, but Figure 3 shows that the resulting TRAC approximations of the SH pressure gradients are remarkably good. These pressure gradients are then responsible for driving the flows.

To that end it is clear the TRAC solution correctly captures the global velocity evolution of the properly resolved SH simulation, during both the ablation and decay phases in response to the 600 s heating pulse. (The 60 s heating pulse simulations show the same fundamental properties). This ensures that TRAC captures the mass and energy exchange that takes place between the chromosphere, TR, and corona correctly, enabling the method to maintain high levels of

accuracy in the coronal temperature and density evolution throughout the full heating and upflow followed by cooling and downflow cycle.

4. Discussion

The difficulty of obtaining adequate spatial resolution in numerical models of the outer solar atmosphere has been a long-standing problem. BC13 demonstrated that lack of adequate spatial resolution during impulsive heating events led to coronal densities that are erroneously small. Johnston et al. (2019) then went on to show that one consequence of these artificially low coronal densities is to (artificially) suppress thermal non-equilibrium (TNE) in coronal loops. Thus, underresolving the TR in numerical simulations has very significant implications for (1) the resultant loop dynamics and (2) any comparisons between model predictions and observations.

Several different approaches have been proposed in order to side-step the need for highly resolved numerical grids and the commensurate very short time steps that are required for numerical stability, yet no fully satisfactory solution is available to date. For example, Lionello et al. (2009) artificially broaden the TR below a fixed cutoff temperature (T_c), while Johnston et al. (2017a, 2017b) dynamically locate the top of an unresolved TR and impose a jump condition that is derived from an integrated form of energy conservation. However, when employed in simulations with coarse spatial resolutions, the latter approach suffers from overestimating the coronal density response to heating while the former, in contrast, underestimates the density.

This Letter has presented a new approach, referred to as the TRAC method, that, for the first time, has successfully removed the influence of underresolving the TR on the coronal density response to heating while retaining remarkable levels of accuracy compared with fully resolved models. The new method combines the basic ideas from the approaches developed previously by Lionello et al. (2009) and Johnston et al. (2017a, 2017b).

We have considered only impulsive heating events where the spatial distribution of the energy deposition is uniform along the loop. However, the TRAC methodology is designed to deal with steep TRs whenever they arise, independent of the nature of the heating. A detailed investigation of different forms of heating (e.g., footpoint heating) will be presented in future work.

TRAC accurately captures the coronal response of properly resolved field-aligned models (e.g., BC13) when employed with spatial resolutions that were up to three orders of magnitude coarser. For example, with TRAC, grid cell widths of order 100 km were sufficient to observe convergence to the fully resolved field-aligned model, which required grid cell widths of order 100 m. The peak density errors were less than 3% and the relaxation of the TR resolution culminated in computation times that were three orders of magnitude faster (e.g., 6 minute run times instead of 3.5 days).

The advantages of this new approach are multiple. For field-aligned hydrodynamic simulations of the coronal response to heating (see, e.g., Reale 2014, for a review), the short computation time means that (1) simulations of coronal heating events can be run quickly, permitting extensive surveys of large parameter spaces (e.g., as done by Froment et al. 2018 to study the occurrence of TNE in coronal loops) to be completed significantly faster at a fraction of the computational cost and

(2) simulations of multiple loop strands (thousands or more) that comprise an entire active region (e.g., experiments seeking to reconcile heating models with the Hi-C observational data), can be performed with relative ease and high accuracy for the coronal emission. However, full numerical resolution is still required to deduce the details of the emission in the TR at temperatures below the adaptive cutoff temperature.

In 3D MHD codes, the method can be included without the need for high spatial resolution in the TR and a corresponding extended computation time, “freeing up” resources to resolve better the current sheets responsible for the heating. Indeed, our results suggest that high levels of accuracy can be obtained with grid cell widths of order 100 km, which is achievable in current 3D MHD simulations.

This research has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 647214). S.J.B. is grateful to the National Science Foundation for supporting this work through CAREER award AGS-1450230. C.D.J. acknowledges support from the International Space Science Institute (ISSI), Bern, Switzerland to the International Team 401 “Observed Multi-Scale Variability of Coronal Loops as a Probe of Coronal Heating.” We also thank Dr Z. Mikić for providing important clarifications on the implementation of the L09 method.

ORCID iDs

C. D. Johnston  <https://orcid.org/0000-0003-4023-9887>
 S. J. Bradshaw  <https://orcid.org/0000-0002-3300-6041>

References

- Antiochos, S. K., & Sturrock, P. A. 1978, *ApJ*, 220, 1137
 Bradshaw, S. J., & Cargill, P. J. 2006, *A&A*, 458, 987
 Bradshaw, S. J., & Cargill, P. J. 2010a, *ApJL*, 710, L39
 Bradshaw, S. J., & Cargill, P. J. 2010b, *ApJ*, 717, 163
 Bradshaw, S. J., & Cargill, P. J. 2013, *ApJ*, 770, 12
 Bradshaw, S. J., & Klimchuk, J. A. 2015, *ApJ*, 811, 129
 Bradshaw, S. J., & Mason, H. E. 2003, *A&A*, 407, 1127
 Cargill, P. J., Bradshaw, S. J., & Klimchuk, J. A. 2012a, *ApJ*, 752, 161
 Cargill, P. J., Bradshaw, S. J., & Klimchuk, J. A. 2012b, *ApJ*, 758, 5
 Cargill, P. J., Warren, H. P., & Bradshaw, S. J. 2015, *RSPTA*, 373, 20140260
 Froment, C., Auchère, F., Mikić, Z., et al. 2018, *ApJ*, 855, 52
 Johnston, C. D., Cargill, P. J., Antolin, P., et al. 2019, *A&A*, submitted
 Johnston, C. D., Hood, A. W., Cargill, P. J., & De Moortel, I. 2017a, *A&A*, 597, A81
 Johnston, C. D., Hood, A. W., Cargill, P. J., & De Moortel, I. 2017b, *A&A*, 605, A8
 Klimchuk, J. A., Patsourakos, S., & Cargill, P. J. 2008, *ApJ*, 682, 1351
 Linker, J. A., Lionello, R., Mikić, Z., & Amari, T. 2001, *JGR*, 106, 25165
 Lionello, R., Linker, J. A., & Mikić, Z. 2009, *ApJ*, 690, 902
 Mikić, Z., Lionello, R., Mok, Y., Linker, J. A., & Winebarger, A. R. 2013, *ApJ*, 773, 94
 Reale, F. 2014, *LRSP*, 11, 4
 Reale, F. 2016, *ApJL*, 826, L20