Research Statement: Jonathan Dursi

I am interested in computational astrophysics broadly — understanding complex astrophysical systems through large-scale parallel multiphysics simulations; carefully examining microphysics both for its own sake and so that it can be included quantitatively in larger-scale models; and developing and analyzing computational techniques for use in astrophysical simulations. I currently study the physics of Type Ia supernovae; the magnetohydrodynamics of galaxy clusters; turbulence in disks; novae; and efficient and accurate computational techniques for studying these phenomenon. Here I outline my plans for these projects for the next few years in the context of the work I have been doing.

Hydrodynamics of Galaxy Clusters

The absence of catastrophic cooling in galaxy clusters has been explained somewhat by recent observations of galaxy clusters, revealing X-ray emission voids of up to 30 kpc in size that have been identified with buoyant, magnetized bubbles. However, the mechanism by which these outflows heat the cluster medium as a whole remains unclear — indeed, how these bubbles continue to exist, much less maintain their observed sharp boundaries, is somewhat mysterious. I have shown that even the modest magnetic fields in the intercluster medium (ICM) may almost automatically be ramped up to dynamically important strengths, affecting both these bubbles and minor mergers. Further investigation will clarify if such effects are directly observable, and will shed light both on the large-scale hydrodynamics of these interesting system, but on the interesting small-scale plasma physics at work in these regimes.

Recent *Chandra* and *XMM-Newton* observations of galaxy cluster cooling flows have revealed X-ray emission voids of up to 30 kpc in size that have been identified with buoyant, magnetized bubbles, presumably inflated by a central Active Galactic Nucleus (AGN). This suggests a way of heating the ICM and preventing a cooling catastrophe. However, while the presence of large amounts of hot gas is suggestive, the process by which the bubble gas heats the ICM is unclear. For example, it is unclear what role thermal diffusion plays, as the ICM magnetic field could potentially be very tangled and all but completely suppress conduction of heat. Turbulence or weak shocks could also play a role. Complicating the issue of bubble/ICM interaction is that how the bubble interface remains sharp is unknown.

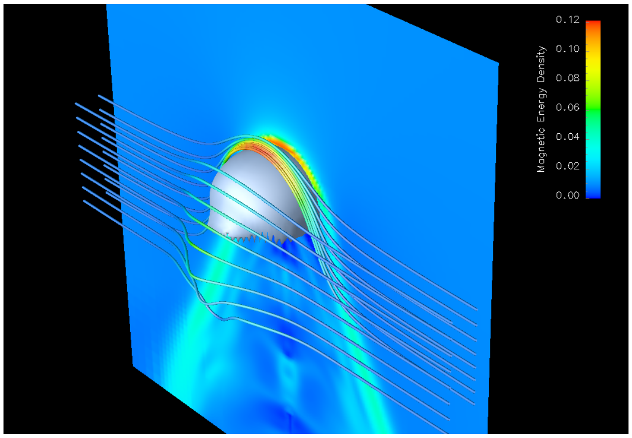

Figure 1: *A projectile in a magnetized medium 'draping' the field over itself, modifying its later dynamics. Movies available here. From Dursi and Pfrommer (2007).*

Recent numerical and analytic work of mine, however, (Dursi, 2007; Dursi and Pfrommer, 2007) has shown that for the case of a small core merging into a magnetized medium — even one quite weakly magnetized, as the $\beta \approx 100$ that one might expect in this context — the projectile 'plows up' fieldlines, building magnetic field strength up to equipartition and beyond, as shown in Fig 1. Further, this magnetic field layer can have significant dynamic effects, from slowing down the projectile to partly stabilizing against instabilities.

These dynamical effects may be quite important, as earlier work of mine done with undergraduate research student K. Robinson at Chicago and P. Ricker at UIUC (Robinson et al., 2004) showed that such a coherent magnetic field may be necessary for the maintenance of these bubbles, as absent any such features, a hydrodynamic bubble will shred itself in one buoyant rise time (*e.g.*, Fig 2), making it very difficult to see how purely hydrodynamic flow could explain both AGN heating and the persistence of such bubbles.

Showing that this works for an over-dense core, however, is not the same as showing that this will also work for an under-dense bubble which is prone to breaking up; however, there is some evidence (Ruszkowski et al., 2007) that it might, and our analysis has confirmed that the timescales could well be sort enough for the draping to successfully provide a protective 'bubble wrap' for these objects, as well. Further, there is the intriguing possibility that these magnetized layers may be directly observable. I plan to continue this work to examine these two questions, also applying the same effect to questions of the ISM.

The work considered above includes only ideal magnetohydrodynamics (MHD), and it is inferred that, because of the magnetic field geometry, thermal conductivity across the drape will be strongly suppressed. However, quantitatively determining this requires self-consistent inclusion of the relevant microphysical thermal and magnetic diffusivities — which are themselves open research topics in this regime. But therein lies opportunity; coupled with the possibility of the direct observation of these drapes, one may be able to constrain what the microphysics must look like, and therefore understand fairly directly the plasma physics of the intercluster medium.

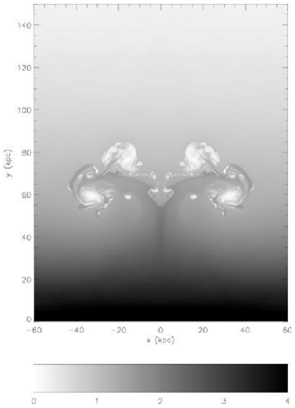

Figure 2: *A bubble self-disrupting as it rises. From Robinson, Dursi, et al., (2004).*

Type Ia Supernovae

Supernova of Type Ia are both fascinating phenomenon in and of themselves, and play an important role in cosmology. However, the underlying mechanism of these explosions remains something of a mystery, and scenarios previously considered 'exotic' may be required to explain recent observations. My past research on these objects has focused on understanding the ignition and burning microphysics that occurs in white dwarfs, including a surprising finding about metallicity dependence; this work is applicable to all plausible scenarios, and is aimed at creating microphysical inputs suitable for use in large-scale simulations. I intend to continue this work, and incorporate it into simulations of large-scale exotic models.

Increasing scrutiny shows Type Ia aren't as uniform as previously thought. What's more, new observations suggest that there may be a 'prompt' population of Type Ia, as well as super-Chandrasekhar mass progenitors, both of which are very difficult to produce in the standard model of supernovae. Understanding these events requires better constraining the scenario which leads to explosion, and a more detailed understanding of the ignition and burning. However, critically examining both ignition/burning and the large scale mechanism simultaneously has been impossible, making distinguishing between scenarios difficult.

My most recent work on Type Ia supernova has focused on the small-scale physics of burning and ignition, with the goal of producing inputs to large-scale simulations. While the overall mechanism remains unclear, it is almost certain that the source of the explosion is the incineration of a carbon-oxygen white dwarf. Thus, questions of how burning would proceed in a degenerate mixture of carbon and oxygen are relevant to all Type Ia scenarios.

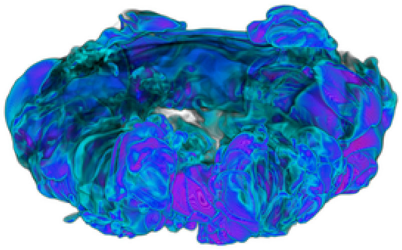

Figure 3: *Early simulations by M. Zingale, in part for Zingale and Dursi (2006), of a rising flame distorted and experiencing shear instabilities; modeled using the novel code employed for the work described here.*

Early work of this type that I was involved with concerned the behaviour of detonations — supersonic combustion that propagates by shock heating. Real detonations are highly multidimensional structures, caused by a cellular instability that corrugates the front, greatly modifying the structure behind the detonation. In particular, the cellular structure can leave material that is only partially burned. This is of great interest for Type Ia supernovae because normally the existence of intermediate-mass (*e.g.*, only partially-burned) elements in the ejecta would disfavor detonations as an explosion mechanism; more detailed understanding of cellular structure allows us to better understand the constraints that observations place on detonation models. This work (Timmes et al., 2000) placed an upper limit on the density at which the cellular instability is likely to produce significant pockets of partially-burned material.

I have also worked extensively on the behaviour of subsonic burning — deflagrations, or flames — which almost certainly are the first stage of burning of these supernovae. The instabilities of these flames are crucial for understanding the speed of the initial burning phase, and has been the focus of most of my work (Dursi et al., 2003; Dursi, 2004; Zingale and Dursi, 2006), explaining the scales on which these flames are stable (see, *e.g.*, Fig. 4).

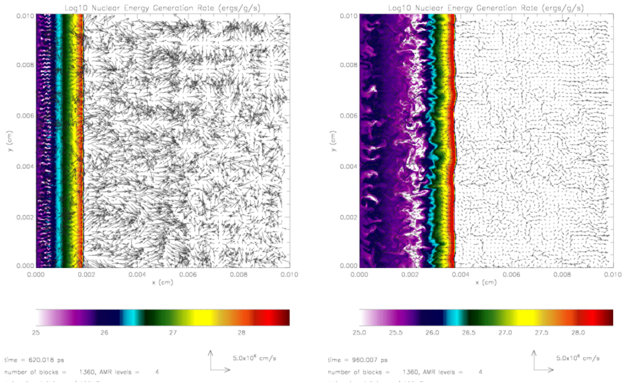

Figure 4: *Flame-turbulence interaction in Type Ia. Note that the flame suppresses small-scale wrinkling at the front, while material behind the front is greatly wrinkled.*

Although much work has now gone into how burning propagates in a white dwarf, the question of how the burning ignites remains an open question – despite the fact that this is a necessary input to any model of a Type Ia, and that differing ignition geometries can strongly affect the final outcome of the explosion. Only recently has attention been focused on ignition physics in a white dwarf.

In Dursi and Timmes (2005), we examined the one-zone ignition timescale for degenerate white dwarf material as temperature, density, and composition vary. This work was initially planned to be a useful but unremarkable tabulation of ignition times, which is a necessary step to more complicated models; however, we found something rather surprising – even quite modest increases in the metallicity of the white dwarf, and in particular the Neon abundances, can significantly decrease the ignition time, significantly raising the probability of a hotspot ignition for more Neon-rich white dwarfs. This, as far as the authors are aware, is the first hint that there could be a difference in the mechanism of ignition between younger and older white dwarfs.

We also examined the ignition of a detonation from a point, and find that a detonation requires a much larger energy input – or coordination on much larger scales – to be successful than is generally understood. Detonations are extremely sensitive to curvature, and even modest curvatures make the propagation of a steady-state detonation impossible. This means detonations at interesting densities in a white dwarf likely require 'matchheads' of up to kilometers in size to launch a detonation. The stage in the explosion of a Type Ia supernovae after ignition is propagation of a burning front and, while detonations may play a role later, at early times the burning propagates as a flame. In this case, where the burning occurs depends sensitively on the balance of the burning, turbulence, and buoyancy; if the buoyancy dominates, then the flame bubble will float harmlessly out of the star, leaving intermediate-mass elements on the surface of the white dwarf but not igniting the star. If the turbulence dominates, the flame surface may be spread over a wide region of the core of the white dwarf, aiding the explosion. In between these regimes, the flame bubble may rise some distance and then become disrupted, essentially beginning the large-scale burning at some finite radius from the center.

With M. Zingale (SUNY SB), in collaboration with researchers at Lawrence Berkeley National Laboratory, I have used a novel computational method along with analytic arguments to examine the balance of these flows on the propagation of flames in Type Ia supernovae (Zingale and Dursi, 2006). Beyond considering when the centre of the star is consumed by the first flames, I have also shown the tendency of a rising flame bubble to fragment (see also Fig. 2), and calculated a characteristic fragmentation scale. This fragmentation scale drops precipitously as degeneracy lifts in the star, providing a rapid cascade of fragmentation – and thus an increase in burning area – at just the densities that an increase in burning rate is needed to match observations. This novel burning model will be examined in future work.

My work so far, then, examines ignition in one-zone (0d) models and early stages of burning in 2d and 3d. My next goal is to extend the work in both directions to build a model which goes from ignition points to early burning phases. This area is currently the biggest unanswered question in Type Ia supernovae models; while the technology has developed to do very sophisticated large-scale simulations of the explosions given some set of initial large-scale burning, they must be initialized with some guess as to what that burning looks like. The goal of my current work is to give a prescription for what that large-scale burning should look like given some progenitor model and the convective turbulence expected from the long-term simmering.

To do this will require extending the consideration of ignition from one-zone models to 1d and higher; this brings in a great deal more physics, and while estimates can be done analytically, precision in the ignition conditions will require very large parameter studies (but of very modest individual computations) to demarcate the 'flammability conditions' for given hotspots. From there, it will be necessary to extend the early burning work done already to include in a more sophisticated way the effects of turbulence and instabilities; the result of these two efforts will be a significant step towards a 'soup-to-nuts' model of the ignition and early burning stages of Type Ia supernovae.

Because the physics of the small-scale ignition does not depend on the large-scale mechanism of the explosion, it will be possible to use this model to critically examine 'exotic' Type Ia mechanisms such as mergers of white dwarf binaries, or pulsational detonation (*e.g.* as in Fig.5).

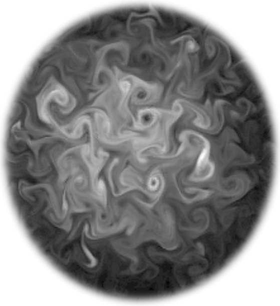

Figure 5: *Results from a simulation of spherically compressed reactive turbulence.*

Turbulence, Mixing, and Instabilities in Disks

Turbulence, shocks, cooling, and fragmentation play important roles in the evolution of disks. The details of transport is determined by the turbulence, and the thermal structure is determined by the shock physics and cooling. Two projects I am currently involved in consider the formation of a turbulent boundary layer at the surface of a disk interacting with a wind; and the interaction of cooling, shock-heating, and accretion in disks around protoneutron stars.

Fragmentation is an important issue in disks, both for angular momentum transport and for planet or star formation. Fragmentation is largely controlled by the balance between irradiation, shock heating, and local cooling process, meaning that getting the cooling rates, shock physics, and dimensionality correct is essential to properly determine limits for the creation of structures.

I am working with R. Rafikov (Princeton) examining questions of gravitational stability of disks; in particular, in situations where a disk is gravitationally unstable (for instance, galactic circumnuclear disks) to examine what sort of statistical properties one would expect; even in very simple cases (*e.g.*, the locally isothermal assumption), very simple quantities such as expected clump mass functions are unknown.

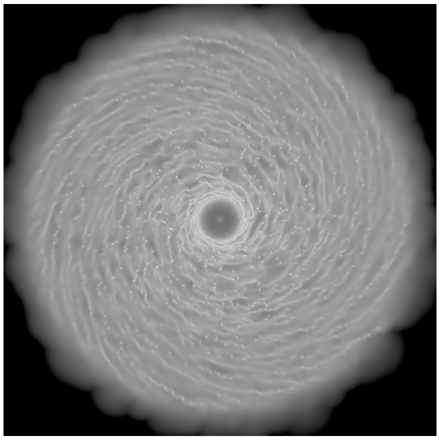

Figure 6: *Gravitational instability in a very thin cold gaseous disk.*

The key here is to obtain sufficient statistics; this can be achived by performing large numbers of individually only moderately-sized calculations (an early example is shown in Fig. 6). We have also determined fairly clever ways of analyzing the resulting reams of data to merge and extract the physically interesting statistical properties. These simulations are being currently performed using the SPH code Gadget; we also expect to perform grid-code calculations to ensure robustness in fragmentation statistics (the effect of numerics of which is a somewhat contentious issue).

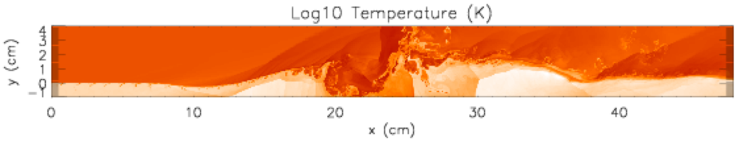

Figure 7: *A hot supersonic wind mixing with the surface of a disk.*

Another project involves wind-disk interactions (eg, Fig 7). If the central object emits a wind, then on the surface of a disk a turbulent boundary layer will be set up, determined by the incoming hot wind and the cooling rate in the disk material. The turbulent boundary layer mediates the ablation of the disk by the wind, and determines the vertical boundary conditions of the disk. I am performing a study of such disk-wind interactions with Chris Thompson (CITA) to understand the interplay of cooling, Kelvin-Helmholtz, and rotation in the formation of the turbulent boundary layer, using extremely large simulations. Mixing properties even in adiabatic supersonic Kelvin-Helmholtz flows are suprisingly poorly understood, and adding even fairly simple cooling physics provides extremely rich dynamics.

Computational Techniques

For both experimentalists and computational scientists, doing new science sometimes means building your own equipment. I have been one of the developers of the FLASH code, a multi-physics, highly parallel, astrophysical fluids code that has won a Gordon Bell prize. But for both groups, doing the next big experiment always requires adding one more tool. In parallel with computational projects using tools already at my disposal, over the next five years I plan to develop 'all-speed' hydrodynamic solvers (already in use in other disciplines) for use in publicly available codes, to make accessible regimes unavailable to the current generation of astrophysical solvers. In testing such a solver, collaboration with experimentalists will be essential, and can generate interesting science in its own right.

My past work on computational techniques for astrophysics has focused on: efficient large-scale parallelization of astrophysical codes; new methods for accurate hydrodynamical simulations; and testing these codes and methods by comparing to semi-analytical results and laboratory experiments. With the FLASH code, I have improved performance of the already robust and efficient PARAMESH AMR package and the individual physics packages, helping FLASH win a Gordon Bell award (Calder et al., 2000), as well as increasing the accuracy of numerical techniques within the code.

Sometimes investigating efficiency of techniques leads to counter-intuitive results, such as our results in (Dursi and Zingale, 2003) which showed that one common technique, 'time subcycling', is of very little use in reducing CPU time in geometrically complex, explicit, AMR calculations. Time subcycling allows the use of fewer timesteps on coarse regions of the grid; however, except for special cases (including many cosmological simulations), these regions make up a small fraction of the zones. Reducing the amount of work on course zones, then, can have only a very modest effect on total compute time, and only at the expense of significant algorithmic overhead.

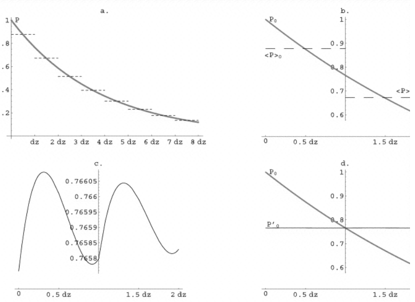

Figure 8: *Modifying the PPM reconstruction to take into account HSE; from Zingale et al (2002)*

For simulations to be useful, calculating them quickly isn't enough; they must also be accurate. With Mike Zingale and (then) undergraduate student John ZuHone, I have developed a set of techniques including modifications to the standard Gudonov-type hydrodynamics solver to increase the accuracy of simulations in stratified atmospheres (Zingale et al., 2002). We showed that this can make the difference between seeing real physics in the stratified atmosphere, or seeing only numerical transients caused by re-settling.

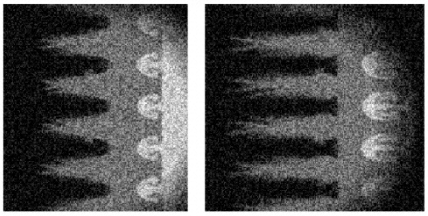

Figure 9: *Simulated radiographs for comparison to a laser experiment; from Calder et al (2002).*

Developing and improving numerical accuracy of an algorithm, however, is only the first level of building confidence in a simulation result. To make sure that simulations accurately reflect reality, contact must continually be made with experimental results. This is as true of established algorithms as they are applied to new problems as it is true of new hydrodynamic methods. I have been closely involved in the comparison of hydrodynamic simulations to laboratory experiments (Calder et al., 2002), and have been involved in the data analysis of one such experiment. This provides not only a way to test simulation codes, but to do new science. An excellent source of astrophysically-relevant laboratory test problems for codes are fluid instabilities. In this case, working with experimentalists not only provides challenging tests for a simulator, but can provide an opportunity to do real, detailed science on instabilities whose complex nonlinear evolution remains an unsolved problem.

An easier, but still valuable, test is to do code-code comparisons, to try to understand the uncertainties in numerical results that come from using different numerical methods (or different implementations of the same methods). I was recently involved in a very large code comparison project (Dimonte et al., 2004) aimed both at understanding the nonlinear evolution of the Rayleigh-Taylor instability and understanding the differences in results between commonly-used codes by application to this problem, and am currently involved in a comparison project examining the decay of supersonic turbulence.

Efficient Parallel Computation

Doing large-scale computation, such as the simulations sketched here, often involves parallel computation, which introduces great possibilities but also great complexities. Many of my contributions to the FLASH code involved improving the efficiency of the parallel computation.

An important stage in this development is the implementation of efficient elliptic solvers. FLASH, which I intend to continue to use as a framework has a multigrid solver which is widely known to be unnecessarily slow. Over the next months — that is, during my remaining time at CITA — I plan to re-write the multigrid solver to greatly improve its performance. This will improve the performance of the FLASH self-gravity module, but will also be an input into other physics solvers.

Hydrodynamic Solvers

Many astrophysical problems involve long periods of relatively quiet, secular evolution followed by a rapid stage of high-speed evolution. Generally, the current generation of multidimensional general-purpose astrophysical hydrodynamic solvers are quite capable at high-speed flows, but timestep restrictions render low speed evolutionary stages inaccessible. Some of my previous work has modestly extended the range of applicability of these methods, but many problems remain inaccessible.

Other disciplines (such as within the combustion research community) have problems with similar natures, and they have developed hydrodynamic solvers capable of dealing with either low-speed flows or even switching smoothly between low- and high-speed flows. In some cases, these methods can be directly used for astrophysical problems; in other cases, fairly significant generalizations must be made because of the highly nonlinear equations of state that occur.

Although there are a variety of methods, they fall largely into two sets of techniques; projection methods and dual-timestepping, representing extending different sorts of techniques into this regime. Projection methods, in particular, use methods similar to those used in other sorts of astrophysical codes, such as for solving for self-gravity, or implicit flux-limited diffusion. It is this class of solvers, then, that I propose to implement and eventually make publicly available.

Validation and Verification

As simulations become more complex, and more important in the understanding of astrophysical phenomenon, it becomes increasingly vital to have a verification and validation strategy for testing a new method. Slowly, more sophisticated approaches are being used in the astrophysical community for verification (ensuring that the model equation are being solved correctly), including the method of manufactured solutions. Validation, however, requires comparison with experimental results.

Low-speed flows in particular lend themselves to comparing simulations with table-top hydrodynamics experiments. For the comparison to be evidence for trusting the astrophysical aspects of the code, the problem being investigated must be a significant part of the target astrophysical problem, but still be simple enough to make a detailed understanding of the results possible. To get experimentalists interested in collaborating, the problem being studied must be interesting in and of itself. These considerations lead naturally to the examination of fluid instabilities, which are important in many astrophysical situations, and generally remain quite poorly understood in the fully nonlinear regime. By pursuing collaborations with experimental groups for investigations of fluid instabilities, one can both test a new solver and do real science with astrophysical application.

References

A. C. Calder, B. C. Curtis, L. J. Dursi, B. Fryxell, G. Henry, P. MacNeice, K. Olson, P. Ricker, R. Rosner, F. X. Timmes, H. Tufo, J. W. Truran, and M. Zingale, editors. *High-Performance Reactive Fluid Flow Simulations Using Adaptive Mesh Refinement on Thousands of Processors*, 2000.

A. C. Calder, B. Fryxell, T. Plewa, R. Rosner, L. J. Dursi, V. G. Weirs, T. Dupont, H. F. Robey, J. O. Kane, B. A. Remington, R. P. Drake, G. Dimonte, M. Zingale, F. X. Timmes, K. Olson, P. Ricker, P. MacNeice, and H. M. Tufo. On Validating an Astrophysical Simulation Code. *ApJSS*, 143:201–229, Nov. 2002.

G. Dimonte, D. L. Youngs, A. Dimits, S. Weber, M. Marinak, S. Wunsch, C. Garasi, A. Robinson, M. J. Andrews, P. Ramaprabhu, A. C. Calder, B. Fryxell, J. Biello, L. Dursi, P. MacNeice, K. Olson, P. Ricker, R. Rosner, F. Timmes, H. Tufo, Y.-N. Young, and M. Zingale. A comparative study of the turbulent Rayleigh-Taylor instability using high-resolution three-dimensional numerical simulations: The Alpha-Group collaboration. *Physics of Fluids*, 16:1668–+, May 2004.

L. J. Dursi. The Linear Instability of Astrophysical Flames in Magnetic Fields. *ApJ*, 606:1039–1056, May 2004.

L. J. Dursi. Bubble-Wrap for Bullets: The Stability Imparted By A Thin Magnetic Layer. *ApJ*, 670:221, Nov. 2007.

L. J. Dursi and C. Pfrommer. Draping of Cluster Magnetic Fields over Bullets and Bubbles – Morphology and Dynamic Effects. *ArXiv e-prints*, 711, Nov. 2007.

L. J. Dursi and F. X. Timmes. Local Ignition in Carbon/Oxygen White Dwarfs – I: One-zone Ignition and Spherical Shock Ignition of Detonations. *ArXiv Astrophysics e-prints*, Oct. 2005.

L. J. Dursi and M. Zingale. Efficiency Gains from Time Refinement on AMR Meshes and Explicit Timestepping. *ArXiv Astrophysics e-prints*, Oct. 2003.

L. J. Dursi, M. Zingale, A. C. Calder, B. Fryxell, F. X. Timmes, N. Vladimirova, R. Rosner, A. Caceres, D. Q. Lamb, K. Olson, P. M. Ricker, K. Riley, A. Siegel, and J. W. Truran. The Response of Model and Astrophysical Thermonuclear Flames to Curvature and Stretch. *ApJ*, 595:955–979, Oct. 2003.

K. Robinson, L. J. Dursi, P. M. Ricker, R. Rosner, A. C. Calder, M. Zingale, J. W. Truran, T. Linde, A. Caceres, B. Fryxell, K. Olson, K. Riley, A. Siegel, and N. Vladimirova. Morphology of Rising Hydrodynamic and Magnetohydrodynamic Bubbles from Numerical Simulations. *ApJ*, 601:621–643, Feb. 2004.

M. Ruszkowski, T. A. Enßlin, M. Brüggen, S. Heinz, and C. Pfrommer. Impact of tangled magnetic fields on fossil radio bubbles. *MNRAS*, pages 400–+, May 2007. doi: 10.1111/j.1365-2966.2007.11801.x.

F. X. Timmes, M. Zingale, K. Olson, B. Fryxell, P. Ricker, A. C. Calder, L. J. Dursi, H. Tufo, P. MacNeice, J. W. Truran, and R. Rosner. On the Cellular Structure of Carbon Detonations. *ApJ*, 543:938–954, Nov. 2000.

M. Zingale and L. J. Dursi. Propagation of the First Flames in Type Ia Supernovae. *ArXiv Astrophysics e-prints*, Oct. 2006.

M. Zingale, L. J. Dursi, J. ZuHone, A. C. Calder, B. Fryxell, T. Plewa, J. W. Truran, A. Caceres, K. Olson, P. M. Ricker, K. Riley, R. Rosner, A. Siegel, F. X. Timmes, and N. Vladimirova. Mapping Initial Hydrostatic Models in Godunov Codes. *ApJSS*, 143: 539–565, Dec. 2002.