FLASH CODE: STUDYING ASTROPHYSICAL THERMONUCLEAR FLASHES

The Center for Astrophysical Thermonuclear Flashes is constructing a new generation of codes designed to study runaway thermonuclear burning on the surface or in the interior of evolved compact stars.

The Center for Astrophysical Thermonuclear Flashes at the University of Chicago—a collaboration between scientists primarily at Chicago and Argonne National Laboratory—studies the longstanding problem of thermonuclear flashes on the surfaces of compact stars (such as neutron stars and white dwarfs) and in the interior of white dwarfs (such as Type Ia supernovae). Our central computational challenge is the breadth of physical phenomena involved. These range from accretion flow onto the surfaces of these compact stars to shear-flow and Rayleigh-Taylor instabilities on the stellar surfaces, nuclear burning ignition under conditions leading to convection, stellar envelope expansion, and the possible creation of a common-envelope binary star system. Many of the physical phenomena we encounter have counterparts in the terrestrial realm (in more extreme versions): convection and turbulence at huge Reynolds and Rayleigh numbers, convective penetration of stable matter, state equations for high-density matter, nuclear processing, radiation hydrodynamics, and interface dynamics (including mixing instabilities and burning front propagation). Perhaps the most spectacular phenomenon with no obvious terrestrial counterpart is the interaction of two binary stars when one of them expands during a nova outburst and swallows its companion (see Figure 1).

The complexity of the physics we must deal with is a common theme at all five ASCI/Alliance Center, and the remarkable physical conditions we encounter distinguishes our center. The energy densities, and many of the physical phenomena, are similar to those dealt with in the US Department of Energy Stockpile Stewardship program. The fully ionized plasmas that ignite under astrophysical conditions are at very high temperatures and densities. The physical problems revolve around nuclear ignition, deflagration or detonation, turbulent mixing, and interface dynamics for complex multicomponent fluids. Our center's task is to develop a code (*Flash code*) that can describe the physics of these astrophysical phenomena and that uses modern software technology to efficiently use available massively parallel computers.

Although the ultimate test of our results is to compare them with astronomical observations

1521-9615/00/$10.00 © 2000 IEEE

ROBERT ROSNER, ALAN CALDER, JONATHAN DURSI, BRUCE FRYXELL, DONALD Q. LAMB, JENS C. NIEMEYER, KEVIN OLSON, PAUL RICKER, FRANK X. TIMMES, JAMES W. TRURAN, HENRY TUFO, YUAN-NAN YOUNG, AND MICHAEL ZINGALE
University of Chicago
EWING LUSK AND RICK STEVENS
Argonne National Laboratory

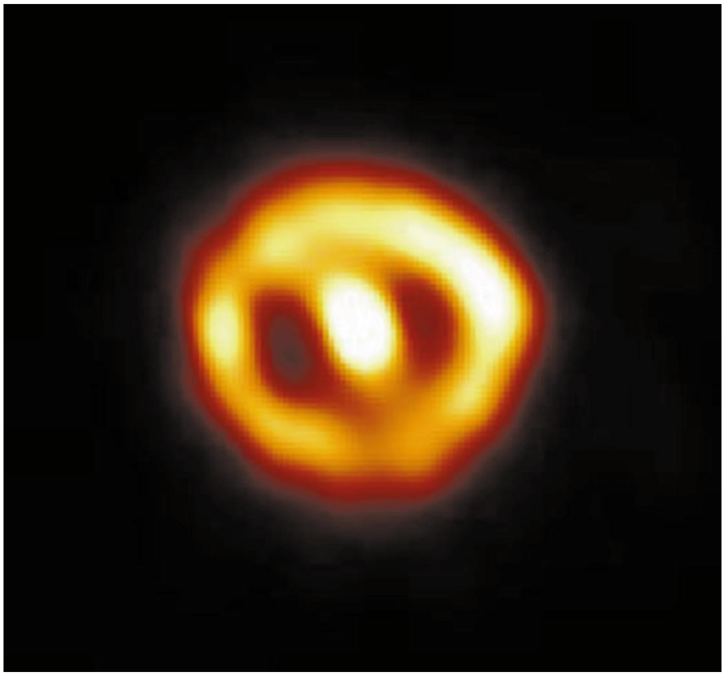

Figure 1. An image of Nova Cygni 1992, a nova in the constellation Cygnus, taken with the Hubble Space Telescope. We can see both the stellar envelope and the expanding shell that resulted from the nova outburst. (Printed with permission from NASA/STScI.)

of the consequences of nuclear flashes on or within compact stars, direct numerical simulations of astrophysical conditions are out of the question, and the direct experimentation familiar to terrestrial physical sciences is simply impossible. Furthermore, the available astrophysical diagnostic tools are extremely limited. Due to these constraints on experimentation and measurement, code validation is key to the astrophysical thermonuclear flash problem.

Because direct numerical simulations are not feasible, researchers have long recognized that astrophysical simulations must ultimately rely on modeling astronomically unresolved small-scale phenomena. Thus, our validation strategy relies on comparing computational results of Flash subsystems (which contain only partial descriptions of the full physics) with model problems for which laboratory verification is possible. These tests range from comparison with various desktop fluid-dynamics experiments to comparison with experiments conducted at national laser facilities (such as Lawrence Livermore National Laboratory and the University of Rochester) and at the pulsed-power facilities at Los Alamos National Laboratory (Pegasus/Atlas) and Sandia National Laboratories (Saturn).

Hydrodynamics: Flash code

Our center has completed the first version of Flash, Flash-1, which addresses the astrophysics problems outlined in the sidebar "Astrophysics background." Visit www.flash.uchicago.edu/flashcode, for details regarding Flash.

For the problem at hand, the code includes these physics:

- *Compressible hydrodynamics.* The current default solver is an explicit higher-order Godunov method based on Phillip Colella and Paul Woodward's piecewise parabolic method,[1] derived in its present form from the Prometheus code.[2]
- *Arbitrary state equations.* Each problem—from astrophysics to validation—requires its own state equation. Typically, we use computationally optimized state equations based on table lookup and interpolation,[3,4] even though in some circumstances far simpler state equations, such as a gamma law, suffice and are available.
- *Arbitrary nuclear-reaction network.* For problems not involving nuclear burning, but only hydrodynamic mixing, the reaction network module can be turned off.[5]
- *External gravity.*

We incorporated these physical effects in Flash-1. This led to a set of equations that govern the motion of compressible matter undergoing nuclear burning in the presence of gravitational stratification, including a continuity equation, a set of advection-diffusion equations that govern the spatial and temporal evolution of each of the nuclear species involved in the reactions, a momentum-conservation equation, and an energy-conservation equation that includes energy input by nuclear burning. We specify a state equation, whose form depends on the particular problem we are treating. In the stellar case, we allow for effects such as electron degeneracy and radiation pressure. At this stage, we assume that the gravitational field is specified (as opposed to computed self-consistently) and that thermal transport is entirely by diffusive processes.

Flash-1

Flash-1 represents a major advance toward a fully flexible code for solving general astrophysical

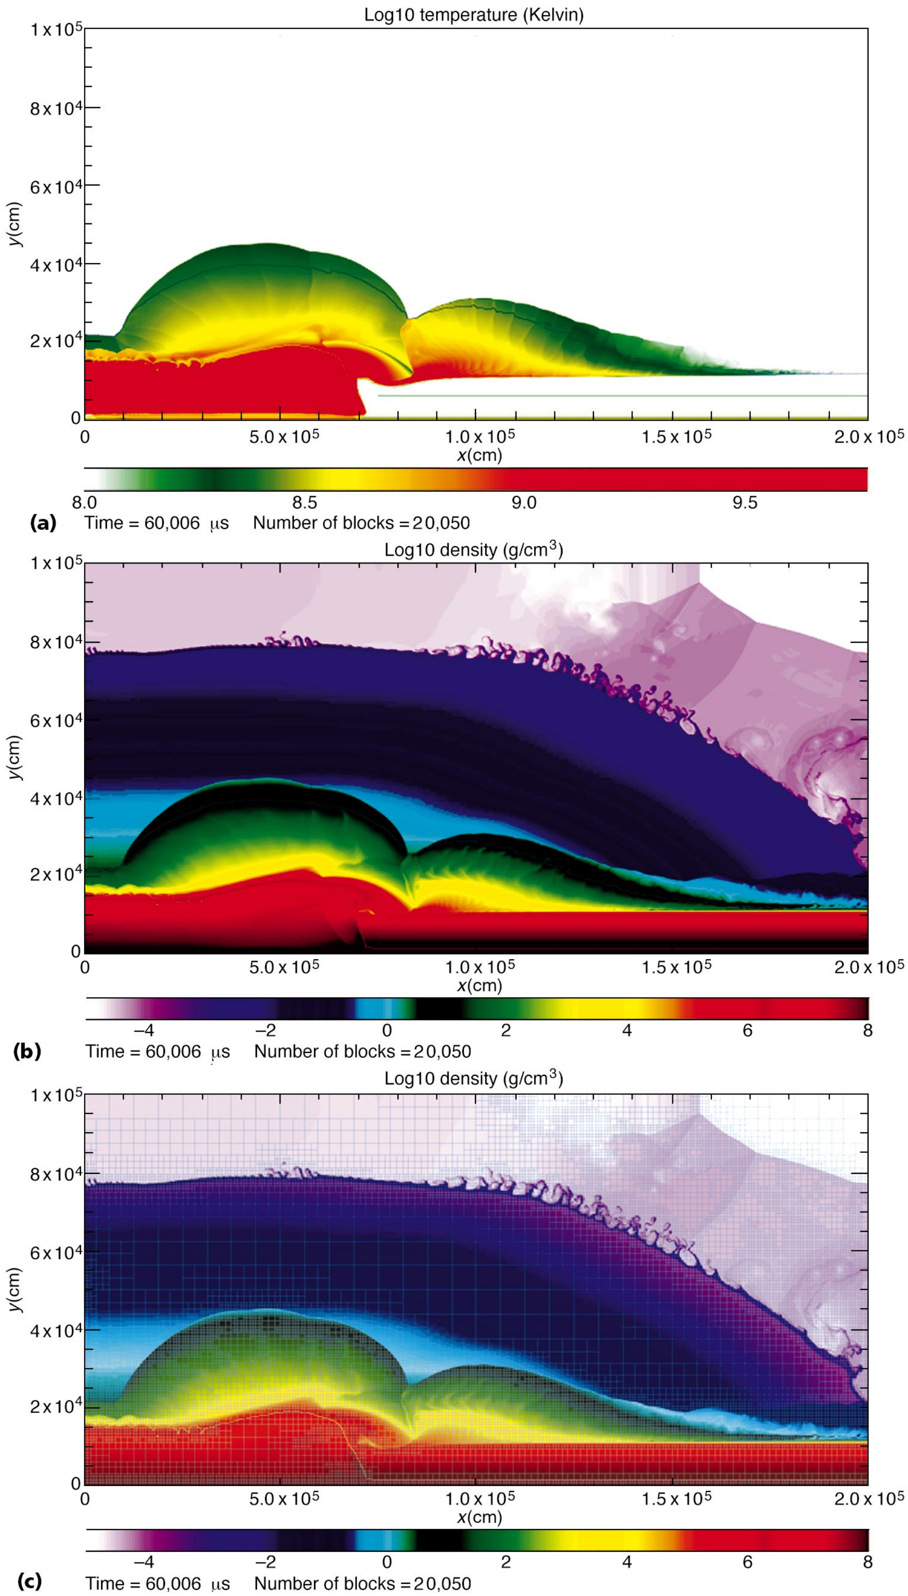

Figure 2: Snapshots of the (a) temperature, (b) density, and (c) adaptive mesh for an X-ray burst on a neutron star's surface. We carried out these computations using Flash-1 on the LANL ASCI platform.

fluid-dynamics problems. Flash-1 is modular and adaptive and operates in parallel computing environments. We designed it to let users configure initial and boundary conditions, change algorithms, and add new physical effects with minimal effort. It uses the Paramesh library to manage a block-structured adaptive grid, placing resolution elements only where we need them most. It also uses the message-passing interface (MPI) library to achieve portability and scalability on a variety of different message-passing parallel computers.[6] To date, we have successfully tested Flash-1 on a variety of Unix-based platforms, including

- SGI systems running IRIX;
- Intel-based systems running Linux (Beowulf systems);
- SGI/Cray T3E running UNICOS;
- the ASCI Blue Mountain machine, built by SGI;
- the ASCI Blue Pacific machine, built by IBM; and
- the ASCI Red machine, built by Intel.

First astrophysical results

It is a reflection on the difficulty of nuclear flash calculations that only recently have researchers progressed beyond Bruce Fryxell and Stanford Woosley's calculation of X-ray bursts.[7] The three panels in Figure 2 show what we can now do with Flash-1 and our results from a 2D cylindrical-coordinates X-ray burst calculation. This particular calculation was performed on the ASCI Nirvana cluster (SGI Origin 2000) at Los Alamos National Laboratory.

Astrophysics background

Runaway thermonuclear burning is relevant to a diverse array of phenomena in astrophysics. Our center focuses on three distinct runaway thermonuclear burning cases: X-ray bursts, novae, and Type Ia supernovae.

X-ray bursts are due to combined hydrogen-helium or pure helium flashes in a shell at the bottom of a thin layer (approximately less than 100 meters) of hydrogen-rich or pure helium material that has accreted onto a neutron star's surface.[1–3] This phenomenon is somewhat simpler than the other phenomena we consider, in that the nuclear energy released per gram of accreted matter is a factor of 20 to 100 less than the gravitational binding energy of the same gram of matter. Consequently, expansion of the envelope does not quench the flash. Rather, the helium and heavier elements in the accreted envelope are incinerated to iron-peak nuclei.

Novae are due to hydrogen flashes in the shell at the bottom of a thin (approximately less than 10^8 cm) layer of hydrogen-rich material that has accreted onto a white dwarf's surface.[4,5] In contrast to the X-ray bursts, the nuclear energy released per gram of accreted matter is a factor of approximately 100 more than the gravitational binding energy of the same gram of matter. As a result, the flash leads to an enormous expansion of the white dwarf's envelope; the envelope engulfs the companion star, forming a common envelope binary. At the same time, the work done against gravity in the expansion of the envelope cools the hydrogen-burning shell and quenches the flash. Steady hydrogen burning then ensues.

Type Ia supernovae are thought to be due to carbon flashes that ignite in the core of a white dwarf whose mass has grown by accretion.[6–9] Neither laminar deflagration nor detonation alone can account for both the abundances of intermediate-mass nuclei and the large expansion velocity of the ejecta that are produced in a Type Ia supernova. Consequently, some Type Ia supernova models invoke a transition from a deflagration wave to a detonation wave. One possibility is that the initial deflagration wave becomes a detonation wave as it travels outward in the star.[8–10] Another possibility is that the initial deflagration wave fails when it reaches the outer part of the star, leading to the white dwarf's recollapse and nearly complete mixing of the nuclear fuel, followed by detonation.[11–17] In either case, much of the white dwarf is incinerated to iron-peak nuclei, and the white dwarf is blown apart.

These phenomena are not only fascinating but also important, because they shed light on other fundamental questions in astrophysics. X-ray bursts tell us about the masses and radii of neutron stars. Classical novae contribute to the abundance of intermediate-mass elements in the galaxy and show how the masses of white dwarfs change with time in close binary systems. Type Ia supernovae contribute to the abundance of intermediate mass and heavy elements in the galaxy. Type Ia supernovae are also important for their crucial role as "standard candles" in determining the Hubble constant and $\Omega_m - \Omega_\Lambda$.[18,19]

Although they seem to be diverse phenomena, X-ray bursts, classical novae, and Type Ia supernovae all involve a close binary star system in which matter from a companion star accretes onto the surface of a compact star (neutron star or white dwarf). All have the ignition of a nuclear fuel under degenerate conditions, followed by the runaway thermonuclear burning through a convective or turbulent flame front (or deflagration wave) or through a shock front (or detonation wave) in common.

Relevant spatial and temporal scales

The three astrophysics problems share a number of common elements that govern their physics and that determine how we simulate these phenomena. As a starting point for discussing the physics, we should consider the typical scale lengths, times, and velocities for nuclear burning under astrophysical circumstances; these make plain the extreme range of spatial and temporal scales this problem spans.

A quick glance at Table A immediately shows that direct numerical simulations of a supernova explosion, for example, would have to span a dynamic range of approximately 10^{15} in time and at least 10^{12} in one spatial dimension. Calculations of such scale are not just impossible to contemplate right now, but it is also dubious whether we will ever achieve them. As long as the stellar interior of a Type Ia supernova, for example, is highly turbulent, adaptive mesh refinement, which researchers might have thought a possible solution, will not help because we will likely carry out refinement almost everywhere within the computational domain, down to at least the Gibson scale (approximately 1 to 10^2 cm).

The underlying physics

We can break down the overall physics problem for thermo-

Table A. Astrophysical scales. T_{burn} is the time it takes for most of the available fuel to be consumed; $\delta_{outer\ scale}$ is the spatial scale of the largest-scale fluid motions; $\delta_{Kolmogorov}$ is the spatial scale at the viscous cutoff.

Scale	Units	X-ray bursts	Type Ia supernova
Time			
$T_{accretion}$	years	10^{-3}	$10^7 – 10^8$
$T_{lightcurve\ rise}$	seconds	<1	10^6
T_{burn}	seconds	3×10^{-3}	~1
Space			
R_{star}	cm	10^6	2×10^8
$\delta_{outer\ scale}$	cm	40	10^7
$\delta_{Kolmogorov}$	cm	10^{-3}	10^{-5}
Speed			
$v_{turbulent}$	cm s^{-1}	5×10^6	10^8

nuclear explosions into several distinct hydrodynamics issues. First, we need to determine the burning front's speed. Detonation and deflagration waves are distinctly different in their computational difficulty. Computing a detonation wave's speed is relatively easy, and in most cases, we can ignore and treat the front's structure as a simple discontinuity. (The exception arises, in some cases, if the front develops a cellular structure,[16,20] which might have an effect on the front's speed.)

In contrast, simulating deflagration waves is much trickier: Unlike detonation, we must always resolve the front's structure to determine a laminar deflagration wave's speed. Unlike novae (for which we can compare the burning front's width with the accreted envelope width[5,21]), the laminar deflagration front's width for Type Ia supernovae can be more than 10 orders of magnitude smaller than the star's radius.[22,23] In principle, we can determine the front's speed in this case from high-resolution 1D simulations[22,23] and can then treat the result as a parameter in the 3D simulations.

Second, we need to establish whether a transition from deflagration to detonation (DDT) occurs. For X-ray bursts, observational evidence points to extremely rapid incineration of the accreted material, and it is a good assumption that burning in this case proceeds through a detonation front. For novae, we know burning is relatively slow (compared to the other two cases we study) and no DDT occurs. For Type Ia supernovae, researchers have long recognized that the transition from a point runaway cannot occur as a prompt detonation. Prompt detonation cannot explain the observed abundances of intermediate mass elements.[12] By considering the observed light curve, we can eliminate a slow deflagration. Therefore, the explosion must involve either a DDT transition or a transition from normal deflagration to turbulent (fast) deflagration.[8–10,15,17]

At a more detailed level, we thus need to understand how chaotic flows in the star affect the propagation of deflagration fronts. Convective instabilities in the burning region and Rayleigh–Taylor and Kelvin–Helmholtz instabilities along the burning front can affect the propagation speed by stretching the flame front and then by introducing small-scale turbulent mixing and energy transport, which might dominate molecular (or microscopic) diffusion processes.[10,15,17] There is no hope (as already mentioned) that we can resolve the deflagration front on a grid that also simulates the behavior on the scale of the entire star, whether or not we are considering nova or supernova evolution. One reasonable approach is to do a high-resolution simulation of a small section of the burning front, to obtain its speed and use the result as a parameter in the full model, combined with a front-tracking method.[24]

The common element of these problems is that we cannot directly compute them with a numerical simulation in the astrophysical context; we will always require models. Therefore, these models must be susceptible to testing; one of the principal aims of our validation effort is to compare such model calculations with direct numerical simulations for laboratory-scale combustion and reaction experiments.

References

1. R.E. Taam, "Nuclear Processes at Neutron Star Surfaces," *Ann. Review of Nuclear Science*, Vol. 35, 1985, pp. 1–23.
2. W.H.G. Lewin, J. van Paradijs, and R.E. Taam, "X-ray Bursts," *Space Science Reviews*, Vol. 62, 1993, pp. 223–389.
3. R.E. Taam et al., "Successive X-ray Bursts from Accreting Neutron Stars," *Astrophysical J.*, Vol. 413, 1993, pp. 324–332.
4. J.W. Truran, "Nuclear Theory of Novae," *Essays in Nuclear Astrophysics*, C.A. Barnes, D.D. Clayton, and D.N. Schramm, eds., Cambridge University Press, Cambridge, Mass. 1982, p. 467.
5. S.A. Glasner, E. Livne, and J.W. Truran, "Reactive Flow in Nova Outbursts," *Astrophysical J.*, Vol. 475, 1997, pp. 754–762.
6. S.E. Woosley and T.A. Weaver, "The Physics of Supernova Explosions," *Ann. Reviews Astronomy and Astrophysics*, Vol. 24, pp. 205–253.
7. K. Nomoto, H. Yamaoka, and T. Shigeyama, "Type I Supernova and Evolution of Interacting Binaries," *Supernovae, Les Houches Session LIV*, S. Bludman, R. Mochkovitch, and J. Zinn-Justin, eds., Elsevier, Amsterdam, 1994.
8. J.C. Niemeyer, 1995, *On the Propagation of Thermonuclear Flames in Type Ia Supernovae,* doctoral dissertation, Tech. Univ. of Munich, 1995.
9. J.C. Niemeyer and W. Hillebrandt, "Turbulent Nuclear Flames in Type Ia Supernovae," *Astrophysical J.*, Vol. 452, 1995, pp. 769–778.
10. J.C. Niemeyer and A.R. Kerstein, "Burning Regimes of Nuclear Flames in SNIa Explosions," *New Astronomy*, Vol. 2, 1997, pp. 239–244.
11. S.I. Blinnikov and A. Khokhlov, "Stage of Spontaneous Flame Propagation in Supernovae," *Soviet Astronomy Letters*, Vol. 13, 1987, p. 364.
12. S.E. Woosley, "Type I Supernovae: Carbon Deflagration and Detonation," *Supernovae*, A.G. Petsckek, ed., Springer Verlag, Berlin, pp. 182–212.
13. A.M. Khokhlov, "Delayed Detonation Model for Type Ia Supernovae," *Astronomy and Astrophysics*, Vol. 245, 1991, pp. 114–128.
14. S.E. Woosley and T.A. Weaver, "Sub-Chandrasekhar Mass Models For Type I Supernovae," *Astrophysical J.*, Vol. 423, 1994, pp. 371–379.
15. A.M. Khokhlov, "Propagation of Turbulent flames in Supernovae," *Astrophysical J.*, Vol. 449, 1995, pp. 695–713.
16. J.R. Boisseau et al., "The Multidimensional Structure of Detonations in Type I Supernovae," *Astrophysical J.*, Vol. 471, 1996, pp. L99–L102.
17. A.M. Khokhlov, E.S. Oran, and J.C. Wheeler, "Deflagration-to-Detonation Transition in Thermonuclear Supernovae," *Astrophysical J.*, Vol. 478, 1997, pp. 678–688.
18. S. Perlmatter et al., "Measurements of Omega and Lambda from 42 High-Redshift Supernovae," *Astrophysical J.*, Vol. 517, No. 2, pp. 565–586.
19. A. Riess et al., "Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant," *AJ*, Vol. 116, No. 3, 1998, pp. 1009–1038.
20. F.A. Williams, *Combustion Theory*, Benjamin-Cummings, Menlo Park, Calif., 1985.
21. S.A. Glasner and E. Livne, "Convective Hydrogen Burning During a Nova Outburst," *Astrophysical J.*, Vol. 445, 1995, pp. L149–L151.
22. F.X. Timmes and S.E. Woosley, "The Conductive Propagation of Nuclear Flames. I. Degenerative C+O and O+Ne+Mg White Dwarfs," *Astrophysical J.*, Vol. 396, 1992, pp. 649–667.
23. F.X. Timmes and S.E. Woosley, " The Conductive Propagation of Nuclear Flames. II. Convectively Bounded Flames in C+O and O+Ne+Mg Cores," *Astrophysical J.*, Vol. 420, 1994, pp. 348–363.
24. M. Reinecke et al., "A New Model for Deflagration Fronts in Reactive Fluids," *Astronomy and Astrophysics*, Vol. 347, 1999, pp. 724–733.

The computational domain is a 2-km region along a neutron star's surface, extending 1 km radially outward. The initial model is a hydrostatic neutron star atmosphere, with a density of 2×10^9 g cm^{-3} at the base, falling off to approximately 10^6 g cm^{-3} over 100 meters. Above this, the density falls off quickly to 10^{-5} g cm^{-3}. The atmosphere and material above it are pure helium. We start the burst by raising the temperature in the domain's lower left corner to 2.5×10^8 K, at which temperature thermonuclear burning of helium commences, creating a detonation. In the vertical direction, the detonation wave propagates down the density gradient, breaking through the surface of the neutron star atmosphere. Above the atmosphere, the shock detaches from the burning front and races ahead. The detonation front moves along the star's surface at approximately 10^9 cm s^{-1}.

In the images in Figure 2, a green line marks a helium abundance of 0.9; below this line, burning has begun to deplete the helium. The dark blue line marks 10 g cm^{-3}, indicating how much the explosion has distorted the neutron star atmosphere. The third panel of Figure 2, showing the superposition of the computational grid on the density field, illustrates this problem's challenging nature. Only with adaptive meshes that can we hope to capture even part of the large dynamic range of spatial scales in this problem.

Code verification and validation

Simulation validation results in the context of the ASCI program revolve around three distinct issues.

- Have we incorporated the correct physics?
- Does the code contain bugs or outright errors?
- Are the computed solutions suitably close to the desired solutions of the mathematical model?

The last question involves both convergence (are the solutions to the discretized equations appropriate approximations to the adopted model equations' solutions?) and a posterior error estimation (are the solutions accurate, or are the computational errors suitably bounded?).

Although it is not trivial to assert with any certainty that a given, highly complex, hydrodynamic code is free from bugs and outright errors, we devised procedures to guard against such difficulties. We constructed a suite of test problems, including the Sod shock-tube problem,[8] the Sedov explosion problem,[8] the Woodward-Colella two-blast-wave problem,[9] an advection problem in which we create a planar density pulse in a uniform pressure region, a double Mach reflection problem, and a wind tunnel flow with a step.

Once we have addressed the problem of outright blunders, we must deal with the remaining two issues, which we only do by comparing simulation results with experimental data. Here we must keep in mind that, even with agreement between measurement and computation, the question remains as to whether the simulation details are trustworthy. Our strategy for choosing validation problems involves four criteria:

- The problem must test significant portions of the full nuclear-flash problem.
- The problem must entail simpler physical problems than the full nuclear-flash problem.
- Tabletop or other laboratory experiments must be available to directly verify code results.
- The problem must be an interesting, forefront physics problem.

Such model problems have several attributes that make them valuable aside from their intrinsic scientific interest. They let us begin the code testing and validation process even in the early versions of Flash. They let us test our methodology for producing modular codes based on reusable components. Finally, they let us explore solution methods or algorithms in simpler situations than the astrophysical case, for which we can more readily understand the physics and for which the diagnostic tools are more powerful (see Figure 3).

Future developments of Flash

Although Flash-1 represents a significant advance in astrophysical hydrodynamics-oriented codes, we want to develop it further. We want to treat additional physical processes and explore (and possibly use) alternate computational algorithms and computational-infrastructure elements that we might need to solve our astrophysics problems. The following descriptions provide a roadmap for our ongoing and future developments and provide insight into what we consider the key bottlenecks in our future computing efforts.

Additional physics

A number of physical effects not yet accounted for in Flash-1 can play important roles in all three of our target astrophysics problems. Future versions of the Flash will take these effects into account.

- *Magnetohydrodynamics* (MHD). Compact evolved stars, such as the ones with which we are concerned, can have dynamically significant magnetic fields, especially at their surface. We need to investigate such magnetic-field effects for both nova and X-ray-burst calculations. Development of an appropriate 3D MHD solver is well underway. Visit www.flash.uchicago.edu for early results.
- *Radiative hydrodynamics*. In all three of our astrophysics problems, radiative transfer effects become important in at least some phases of the flash evolution. As a result, Flash's treatment of radiative energy transport, which is currently restricted to the simple diffusion approximation, will need to include at minimum flux-limited diffusion.
- *Self-gravity*. Depending on the problem, we might need to calculate the gravitational field as part of the problem's solution (for example, in the case of supernova evolution). This requires adding a Poisson solver, which we currently plan to implement using our newly developed multigrid solver for Flash-1.
- *Rotation*. Rotation strongly affects fluid motions whose characteristic time scale is comparable to, or longer than, the rotation period. Such effects might be important in our Flash calculations but have received relatively little attention in the literature.

Figure 3. A Rayleigh-Taylor instability validation problem, the instability of a heavy fluid supported by a relatively light fluid. In this calculation, we look at Rayleigh-Taylor in a weakly compressible (Boussinesq) fluid and study the flows leading to chaotic mixing. We obtained the image from a 3D pseudospectral calculation. It shows a vertical slice through the computational domain, illustrating the complex plume structure of the mixing region. The calculation uses 256 × 256 × 512 modes and was carried out at the Pittsburgh Supercomputing Center under ASCI sponsorship.

Algorithms

It is widely understood, although often forgotten in the telling, that algorithmic improvements can have impacts on computing effectiveness that are at least as profound as the changes resulting from improvements in computing hardware. We have identified three algorithmic issues that are particular to our Flash problem and on which we are focusing: simplified hydro solvers, unstructured grids, and discontinuous Galerkin techniques

In certain astrophysical regimes, fluid motions are very subsonic. Gravitational stratification might also be weak; an example is convection near or at the center of an evolved star. In such cases, filtering out sound waves (the anelastic approximation) and, if permissible, ignoring gravitational stratification (leading to the Boussinesq approximation) save considerable computational effort. We are pursuing two complementary avenues to address these simplifications. First, we plan to implement a spectral element hydro solver[10] (now used at ANL) in the Paramesh framework of the current Flash. Second, our computational physicists are developing a semi-implicit hydro module for compressible hydrodynamics.

In astrophysical applications, we rarely encounter problems with geometrically complex boundaries. This together with the relative simplicity and effectiveness of coding on structured meshes, explains the general lack of interest in unstructured meshes among astrophysicists. However, we might more easily or accurately solve certain fluid-dynamics problems for a given effort level using unstructured meshes. One family of examples is the class of converging-flow problems (implosions). To better understand the computational issues, we are examining the use of unstructured meshes, based both on the SUMMA3d and RPI architectures.

In some instances (such as when we use unstructured grids), dealing with a very compact stencil for the spatial differential operators is preferable. We are conducting a comparative study of one such family of techniques using compact stencils—namely, discontinuous Galerkin methods—using the Rayleigh-Taylor validation problem as our problem test bed. This study also compares results obtained with discontinuous Galerkin methods on structured meshes with results obtained with Flash-1.

Large-scale computing research

Our ANL colleagues, in collaboration with colleagues at Chicago and at the DP national laboratories, are pursuing a number of computer science research issues that will improve our ability to carry out large-scale computations:

- Performance diagnostics for large-scale parallelized codes is a challenging task that the ANL Jumpshot development effort[11] hopes to simplify. We have used this tool to take a detailed look at Flash-1's performance on Blue Pacific. We plan to use a fully portable implementation on all the available massively parallel computing platforms.
- A major limitation of currently available software for parallel I/O is the absence of uniformly implemented standards on available massively parallel platforms. Although our planned long-term solution relies on the HDF5 standard, in the short term, we are using MPI-IO, developed at ANL. Therefore, we are relying strongly on the current effort to produce a fully portable MPI-IO version.
- A major bottleneck of present large-scale computing efforts is the limited effectiveness of end-to-end data transmission (such as from the computing sites to our local desktops) for very large data sets. ANL scientists helping our center resolve this issue, principally through the data grid initiative. We hope to extend the substantial performance improvements (with transmission speeds of 40 Mbytes) now demonstrated for data movement between ANL and Lawrence Berkeley National Laboratory over ESnet to data movement between the labs and our center.
- In an ideal Flash working environment, the astrophysicist sitting at the desktop machine would be able to run code without worrying about exactly where the code is actually running or where the resulting data is stored. To reach this goal, we have prototyped a Globus-based problem-solving environment for Flash-1.
- Computational physicists and computer scientists at our center are exploring advances in program architectures, including the use of object-oriented code frameworks for structured-mesh calculations (in particular, Samrai [see www.llnl.gov/casc/Samrai]) and incorporation of code-module interface standards developed as part of the common component architecture and equation solver interface efforts. These developments will be key to Flash's future evolution.
- We cannot easily visualize or analyze the huge data sets we are now beginning to deal with available commercial or public domain software. Research at ANL is particularly aimed at such problems and is exploring techniques such as adaptive multiresolution visualization (based on use of the Alice-based memory snooper) and parallel data reduction (using a distributed octree data structure), and the use of high-performance visualization hardware tools, such as the CAVE/Immersadesk and the ActiveMural.

Most importantly, Flash simulations must give the correct answer. This is difficult because astrophysicists are fundamentally constrained to watch distant events without any ability to control them or to diagnose them by in situ measurement. So, our confidence in our computational results hinges crucially on the ability to validate calculations using appropriate proxy measurement or code validation using appropriate laboratory models.

Acknowledgments

We thank our many colleagues at Chicago, Argonne, and Rensselaer Polytechnic Institute for their contributions to and general enthusiasm for our project. The Flash Center is supported under the DOE ASCI/Alliances Program at the University of Chicago.

References

1. P. Colella and P.R. Woodward, "The Peicewise-Parabolic Method (PPM) for Gas-Dynamical Simulations," *J. Computational Physics*, Vol. 54, 1984, pp. 174–201.
2. B. Fryxell, E. Müller, and W.D. Arnett, *Hydrodynamics and Nuclear Burning*, Max-Planck-Institut für Astrophysik, Report 449, 1989.
3. F.X. Timmes and W.D. Arnett, "The Accuracy, Consistency, and Speed of Five Equations of State for Stellar Hydrodynamics," *Astrophysical J. Supplement Series*, Vol. 125, 1999, p. 294.
4. F.X. Timmes and F.D. Swesty, "The Accuracy, Consistency, and Speed of an Elctron-Positron Equation of State Based on Table Interpolation of the Helmholtz Free Energy," to be published in *Astrophysical J. Supplement Series*, Feb. 2000.
5. F.X. Timmes, "Integration of Nuclear Reaction Networks for Stellar Hydrodynamics," *Astrophysical J. Supplement Series*, Vol. 124, 1999, pp. 241–263.
6. W. Gropp, E. Lusk, and A. Skjellum, *Using MPI: Portable Parallel Programming with the Message-Passing Interface*, 2nd ed., MIT Press, Cambridge, Mass., 1999.
7. B.A. Fryxell and S.E. Woosley, "A Two-Dimensional Model for Gamma-ray Bursts," *Astrophysical J.*, Vol. 258, 1982, pp. 733–739.
8. L.I. Sedov, *Similarity and Dimensional Methods in Mechanics*, Academic Press, New York, 1959.
9. P. Woodward and P. Colella, "The Numerical Simulation of Two-Dimensional Fluid Flow with Strong Shocks," *J. Computational Physics*, Vol. 54, 1984, pp. 115–173.
10. P.F. Fischer, N.I. Miller, and H.M. Tufo, "An Overlapping Schwarz Method for Spectral Element Simulation of Three-Dimensional Incompressible Flows," *Parallel Solution of Partial Differential Equations*, P. Bjorstad and M. Luskin, eds., Springer-Verlag, New York, 1999, pp. 159–180.
11. O. Zaki et al., "Toward Scalable Performance Visualization Jumpshop," *Int. J. High Performance Computing Applications*, Vol. 13, No. 3, 1999, pp. 277–288.

Robert Rosner is the William E. Wrather Distinguished Service Professor in the Department of Astronomy and Astrophysics, the Enrico Fermi Institute, and the College at the University of Chicago and is the director of the Flash Center. His interests range from solar physics and stellar activity to plasma astrophysics and fluid dynamics. Contact him at the Dept. of Astronomy and Astrophysics, Univ. of Chicago, 5640 S. Ellis Ave., Chicago, IL 60637; r-rosner@uchicago.edu.

Alan Calder, **Jonathan Dursi**, **Bruce Fryxell**, **Donald Q. Lamb**, **Jens C. Niemeyer**, **Kevin Olson**, **Paul Ricker**, **Frank X. Timmes**, **James W. Truran**, **Henry Tufo**, **Yuan-Nan Young**, and **Michael Zingale** are members of the astrophysics group and Flash team at Chicago who have been primarily responsible for building Flash-1.

Ewing Lusk and **Rick Stevens** have been guiding Argonne National Laboratory's contributions to the Flash project.