Supernova Simulation with CHIMERA

O. E. Bronson Messer
National Center for Computational Sciences
Oak Ridge National Laboratory

Stephen W. Bruenn
Department of Physics
Florida Atlantic University

W. Raphael Hix & Anthony Mezzacappa
Physics Division
Oak Ridge National Laboratory

John Blondin
Department of Physics
North Carolina State University

Nathan Wichmann
Cray Inc.

Abstract

CHIMERA is a multi-dimensional radiation hydrodynamics code designed to study core-collapse supernovae. The code is made up of three essentially independent parts: hydrodynamics, nuclear burning, and a neutrino transport solver combined within an operator-split approach. The multi-physics nature of the problem, and the specific implementation of that physics in CHIMERA, provide a rather straightforward path to effective use of multi-core platforms in the near future.
KEYWORDS: astrophysical hydrodynamics, numerical methods, radiative transfer, supernovae

1 Introduction

Core-collapse supernovae are among the most energetic events in the Universe, releasing 10^{53} erg of energy on timescales of a few tens of seconds. They produce and disseminate many of the elements heavier than helium, making life as we know it possible. They mark the birth of neutron stars and black holes. In recent years, it has become apparent that core-collapse supernovae from massive progenitors are associated with long gamma-ray bursts (Matheson et al., 2003; Galama et al., 1998; Gehrels et al., 2005)

As the name suggests, core-collapse supernovae are initiated by the collapse of the iron cores of massive stars at the ends of their lives. The collapse proceeds to ultrahigh densities, in excess of the densities of nucleons in the nucleus of an atom (supernuclear densities). The inner core becomes incompressible under these extremes, bounces, and, acting like a piston, launches a shock wave into the outer stellar core. This shock wave will ultimately propagate through the stellar layers beyond the core and completely disrupt the star in an explosion. However, in all realistic simulations to date, the shock stalls in the outer core, losing energy as it plows through the still infalling material. Exactly how the shock is revived is unknown. This is the central question in core-collapse supernova theory. Answering this question is made all the more difficult by

the daunting amount of physics inherent in the phenomenon. Core-collapse supernovae are incredibly complex events, requiring input from all the major fields of modern physics – nuclear physics, particle physics, relativity, solid-state physics – and modern computational methods to describe them adequately. For all these reasons, the ultimate core-collapse supernova mechanism remains unknown despite four decades of computational effort.

Evidence has accumulated indicating that multidimensional effects play an important and perhaps essential role in the mechanism. On the observational side, spectropolarimetry, the large average pulsar velocities, and the morphology of highly resolved images of SN 1987A all suggest that anisotropy develops very early on in the explosion [e.g., see Arnett et al. (1987) and McCray (1993) for reviews and references]. On the theoretical side, analyses of immediate post-bounce core profiles given by computer simulations show that a variety of fluid instabilities are present and may play a role in the explosion mechanism [e.g., see Buras et al. (2006), for a review]. In particular, multi-dimensional numerical simulations have shown that convective overturn in the neutrino-heated region behind the stalled shock may be important for the success of the neutrino-driven mechanism, as it helps transport hot gas from the neutrino-heating region directly to the shock, while downflows simultaneously carry cold, accreted matter to the layer of strongest neutrino heating where a part of this gas readily absorbs more energy from the neutrinos. These simulations have also revealed that a non-radial, low-mode standing accretion shock instability (SASI) may also grow, given time, via the propagation of sound waves (Blondin & Mezzacappa, 2007). This low-mode distortion of the shock may be at the root of some of the above mentioned supernova observables.

The complexity of the supernova mechanism precludes a purely analytic investigation, requiring, instead, realistic numerical simulations. This presents great technical challenges. A typical supernova explosion energy is 10^51 ergs, or 1B (a unit of energy, the bethe, which honors Hans Bethe, who spent more than a decade contributing to core collapse supernova theory), and must be regarded as marginal, being of the same order as the gravitational binding energy of the envelope of the progenitor ejected. On the other hand, 100 times this energy resides in the internal energy of the immediate post collapsed core, and the near negative of this resides in the form of gravitational binding energy. Thus, simulations must be energy conserving to high accuracy if we are to take their outcomes seriously. Ultimately, 300 B in energy is released by the core in neutrinos of all flavors, and their interaction with the stellar core and mantle will either power the explosion itself or play a ma jor role in the explosion dynamics. An inaccurate treatment of neutrino transport can qualitatively change the results of a simulation. Since neutrinos can originate deep within the core, where neutrino mean free paths are small compared with other relevant length scales, and propagate out to regions where the reverse is true, the transport scheme must be accurate in both regimes plus the all-important intermediate regime where the critical neutrino energy deposition occurs. Neutrinos interact with matter in a variety of energy-dependent ways, and this demands that both the neutrino transport and the interactions receive a full spectral implementation, rather than having the neutrino spectrum prescribed. The angular distribution of the neutrinos is also important to compute accurately. In particular, it affects the neutrino heating, and the latter is primarily determined in a region where the angular distribution can neither be assumed to be isotropic nor radially free streaming. Supernova simulations must be carried out in two, and preferably three, spatial dimensions for the reasons mentioned . The nuclear abundances should be evolved in regions where nuclear statistical equilibrium (NSE) cannot be maintained. This will enable the potentially observable products of nucleosynthesis to be followed and the energy released by nuclear burning to be fed back into the computation of the explosion dynamics. While the energy released is expected to be rather small, it could be locally significant and have an influence on the dynamics if all other factors give rise to a very marginal outcome. Finally, general relativistic effects must be incorporated, as they influence the size of the neutrino heated region, the rate of matter advection through this region, and the neutrino luminosities and RMS energies (Bruenn et al., 2001). They can profoundly affect the dynamics.

2 CHIMERA architecture

CHIMERA can well be described as a "chimera" of three, separate, rather mature codes. The codes are tightly coupled through a central database-type data

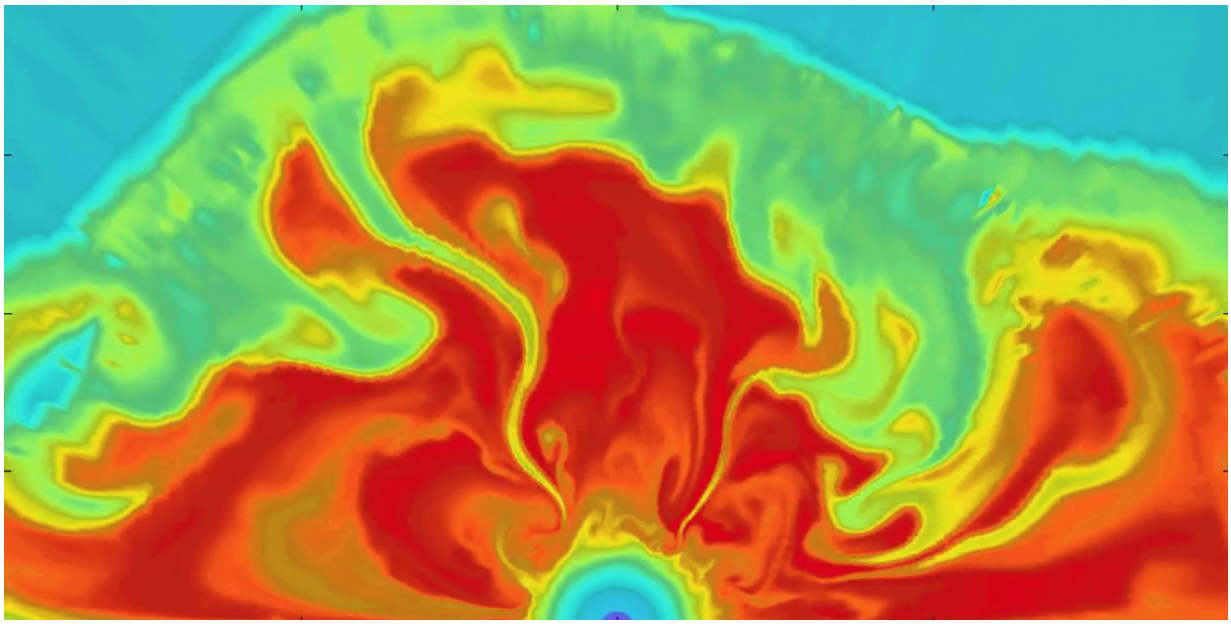

Figure 1: Visualization of the stellar matter entropy in a 2D CHIMERA simulation. The snapshot is from the early evolution of an explosion of a 11.2 solar mass progenitor star. The small blue (cool) hemisphere at the center is the nascent neutron star.

structure (cf. the FLASH code), where each code module "checks out" copies of global variables, uses or evolves them, and then "checks them back in." The primary code modules are used to evolve the stellar gas hydrodynamics (VH1/MVH3), the ray-by-ray neutrino transport (MGFLD-TRANS), and the nuclear kinetics (XNET). These three "heads" are augmented by a sophisticated equation of state for nuclear matter (e.g. Lattimer & Swesty (1991)) and a self-gravity solver capable of an approximation to general-relativistic gravity(specifically, a spectral Poisson solver is used to determine the gravitational field (Müller & Steinmetz, 1995) with general-relativistic corrections to the spherical component (Marek et al., 2006)).

2.1 Hydrodynamics

The hydrodynamics module in CHIMERA is a modified version of the PPM code VH-1, which has been widely used in astrophysical fluid dynamics simulations and as an important benchmark code for a variety of platforms. VH-1/MVH3 is a Lagrangian remap implementation of the Piecewise Parabolic Method (PPM) (Collela & Woodward, 1984). Being third order in space (for equal zoning) and second order in time, the code is well suited for resolving shocks, composition discontinuities, etc. with modest grid requirements. Redshift and time dilation corrections are included in both the hydrodynamics and neutrino transport (to be described later). A moving radial grid option, where the radial grid follows the average radial motion of the fluid, makes it possible for the core infall phase to be followed with good resolution. The code module has recently been updated with a new domain decomposition scheme more amenable to three-dimensional simulations in spherical polar coordinates (a necessary step for performing three-dimensional simulations within the ray-by-ray approach to be described below).

2.2 Neutrino Transport

Ideally, neutrino transport should be implemented with full multi-D Boltzmann transport. This important effort is being made but will be very computationally expensive. We compromise by implementing a ray-by-ray-plus approximation [cf. Buras et al. (2006)] for neutrino transport, whereby the lateral effects of neutrinos such as lateral pressure gradients (in optically thick conditions), neutrino advection, and velocity corrections are taken into account,

but transport is performed only in the radial direction. Transport is computed by means of multigroup flux- limited diffusion with a flux limiter that has been tuned to reproduce Boltzmann transport results to within a few percent (Liebendörfer et al., 2004) . All $O(v/c)$ observer corrections have been included. The transport solver is fully implicit and solves for four neutrino flavors simultaneously (i.e., ν_e's, $\bar{\nu}_e$'s, ν_μ's and ν_τ 's (collectively ν_x's), and $\bar{\nu}_\mu$'s and $\bar{\nu}_\tau$'s (collectively $\bar{\nu}_x$'s)), allowing for neutrino neutrino scattering and pair-exchange, and different ν and $\bar{\nu}$ opacities. The PPM technology has been directly applied to both the spatial and energy advection of neutrinos in both the radial and lateral directions. The neutrino opacities employed for the simulations are the standard ones described in Bruenn (1985), with the isoenergetic scattering of nucleons replaced by the more exact formalism of Reddy et al. (1998), which includes nucleon blocking, recoil, and relativistic effects, and with the addition of nucleon-nucleon bremsstrahlung (Hannestad & Raffelt, 1998) with the kernel reduced by a factor of five in accordance with the results of Hanhart et al. (2001).

2.3 Nuclear Kinetics

The nuclear composition in the non-NSE regions of these models is evolved by the thermonuclear reaction network of Hix & Thielemann (1999). This is a fully implicit general purpose reaction network; however,currently we have implemented only a so-called α-network, i.e. only reactions linking the 14 alpha nuclei from ^{4}He to ^{60}Zn are used. Data for these reactions is drawn from the REACLIB compilations (Rauscher et al., 1996). The nucleons have only very small abundances at any time and are included to make the NSEnon-NSE transition smoother. The iron-like nucleus is included to conserve charge in a freezeout occurring with an electron fraction below 0.5 [cf. Kifonidis et al. (2003)].

3 Module Performance

3.1 Hydrodynamics Scaling

Although the hydrodynamics is a very small part of the overall cost of a supernova simulation (e.g. for a fiducial resolution of 600 radial zones and 200×100 angular zones, a single hydrodynamic timestep, including communication, on 5000 cores takes less than a third of second, whereas the nuclear network solve takes >10 seconds and the neutrino transport solve requires>3 seconds), the number of desirable zones in a high-resolution simulation will dictate the use of >10,000 to 20,000 or more MPI processes. Therefore, excellent scaling at lower processor counts is essential for future success.

The new domain decomposition for MVH3 is shown in Figure 2. Hydrodynamic sweeps are made on "pencils" along one direction of a logically Cartesian mesh. Then, a data transpose is performed to switch the sense of the sweeps to one of the orthogonal directions, followed by a sweep. This procedure can be interleaved in various ways within the operator-split scheme, but a canonical hydro timestep would have sweeps like X-Y-Z-Z-Y-X, i.e. a sweep in the "X-direction" (or, e.g. radius), followed by a Y sweep, followed by a Z sweep, followed by a reverse of that sequence. This decomposition is necessary for the ray-by-ray neutrino transport, as it allows a single "ray" to be resident on a processor at some point in a timestep. This makes the neutrino transport solve a wholly local computation, requiring no communication.

We have performed a series of benchmark hydro-only simulations at a set resolution of 600 radial by 200 azimuthal zones by 100 longitudinal zones. Results of these benchmarks are shown in Figure 3. Communication time is broken down into time for "gets" and "puts" for a each sweep step.

MVH3 starts to exhibit reduced scaling above 2500 cores and after 5000 cores the lack of local work per process serves to effectively halt scaling. This can be seen in the breakdown of communication and computation. While computation continues to scale, the communication does not and accounts for almost 50% of the time at 5000 cores. Although the local work per timestep will increase many orders of magnitude in the full application, making the hydro solves themselves as scalable as possible is a design goal.

We have investigated a new implementation of a modified version of MVH3 on the X1E using Co-Array Fortran (CAF) to do "puts" into position for the next sweep. The parallel loops are blocked into "pencils" and puts are done as pencils are completed. This modification does three things:

- It eliminates the "get" stage.
- As the puts are done as pencils the communication is spread out in time much more as opposed to being focused in an single all-to-

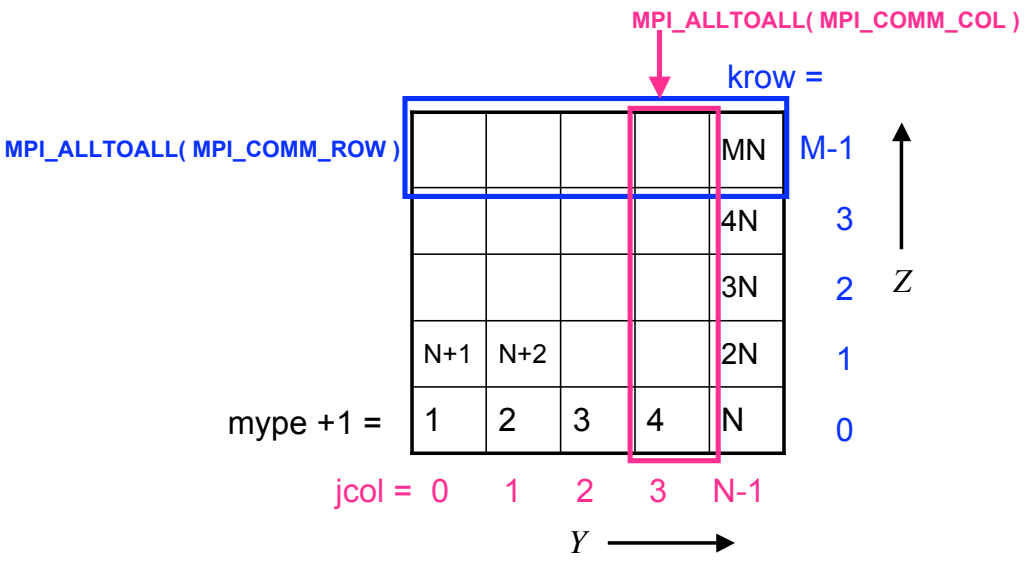

Figure 2: Schematic view of MVH3 domain decomposition.

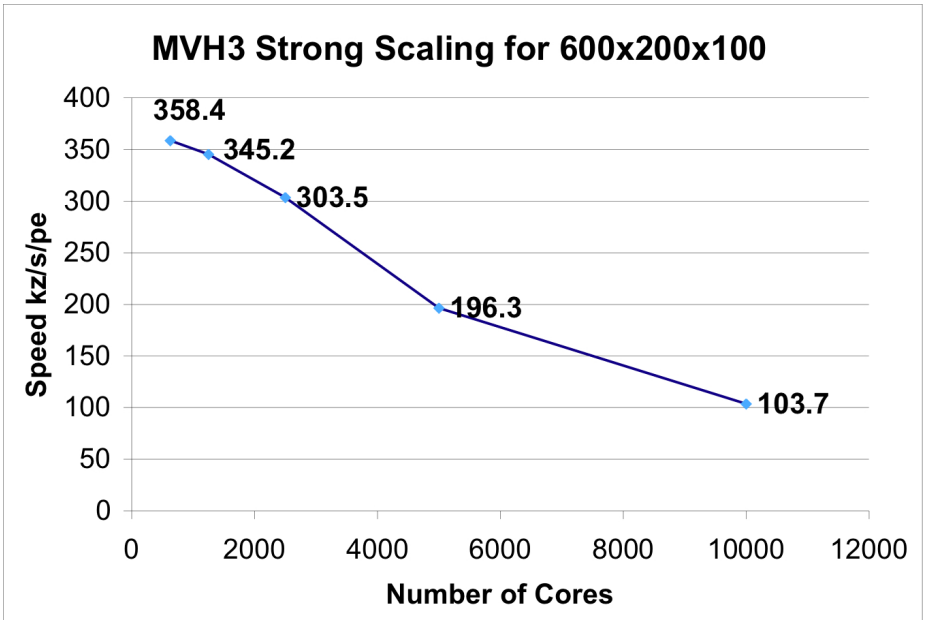

Figure 3: MVH3 strong scaling for a typical resolution for supernova simulation.

all.

- The puts can be asynchronous, meaning virtually all of the latency can be hidden.

This has been done on the X1E as we had access to a CAF compiler for more rapid prototyping. In the future, we hope to convert the CAF version to use the shmem library on the XT platform. This is future work not entirely limited by our time and interest, however: the XT3 was not designed to do one-sided communication. However, Baker is designed for one-sided communication in many ways, and shmem performance should be quite good. In any event, this should be a rather robust use of the model, and shmem developers should be able to use it as an effective foil for development.

3.2 XNET Single-processor Performance

As the calculation of the nuclear kinetics will become the dominant computational load as we move to a 150-isotope network while keeping the neutrino transport in the MGFLD approximation, the single-processor performance of XNET becomes an important consideration. We have profiled a typical supernova nucleosynthesis problem using XNET and find that the building of the Jacobian for the implicit solve takes roughly half the 5 seconds necessary to evolve a single hydro zone. The other half of the time is in the solve itself. This solve, since it lies at the bifurcation point between the sizes where dense and sparse solvers are found to work best, is simply performed with a single LAPACK driver call. Gratifyingly, we find that the entire module, owning to reasonable construction of the Jacobian and good vendor-supplied LAPACK and BLAS performance, achieves a little better than 50% of peak on the XT4.

4 Conclusion

We have developed a numerical code coupling multi-dimensional hydrodynamics, a nuclear reaction network, and spectral neutrino transport in a "ray-by-ray-plus" approximation to simulate core collapse supernovae from the infall epoch to ≈1 sec post bounce. We have already performed 2D simulations with the code, obtaining explosions for a wide variety of progenitor masses and incorporated physics. High-entropy plumes separated by lower-entropy downflows are evident by 100 ms post bounce for all models. The plumes begin to merge, so that by 200 ms two or three large bubbles remain and begin to distort the shock. Bipolar oscillations

of the shock begin to become evident at this time as well, which we attribute to the SASI. Most interesting is the synergistic interplay between the reduced ram pressure, the energy released by the burning of the shocked oxygen-rich material, and neutrino heating that ensues once the oxygen-rich layers penetrate the shock. All three ingredients appear to be essential and result in the shock being pushed out into the unburnt material and an explosion. The results reported here are very promising in that many supernova observables may be reproduced, but they need to be viewed with caution. The simulations need to repeated with GR incorporated into the code, and in 3 dimensions, preferably with the use of a singularity-free grid.

5 About the Authors

Bronson Messer is an R&D Staff Member in the Scientific Computing Group of the National Center for Computational Sciences at Oak Ridge National Laboratory. He spends most of his time being the group's "pet astrophysicist." Bronson is also a matrixed member of the Theoretical Astrophysics Group in the Physics Division at ORNL. He can be reached at bronson-at-ornl-dot-gov.

Steve Bruenn is professor of Physics at Florida Atlantic University. He is the author of MGFLD-TRANS and the primary designer and code architect for CHIMERA. His 1985 paper on the initial results from MGFLD-TRANS is one of the most commonly cited works in supernova theory, even now. He can be reached at bruenn-at-fau-dot-edu.

Raph Hix is an R&D Staff Member in the Theoretical Astrophysics Group in the Physics Division at ORNL and an adjunct associate professor of physics at the University of Tennessee, Knoxville. He is the author of the nuclear network code XNET. He can be reached at raph-at-ornl-dot-gov.

Tony Mezzacappa is a Corporate Fellow at ORNL and head of the Theoretical Astrophysics Group. He is PI of the INCITE project, "Multidimensional Simulations of Core-Collapse Supernovae," representing the single largest award on the NLCF XT3/4 system in FY2007. He can be reached at mezzacappaa(note the extra 'a')-at-ornl-dot-gov.

John Blondin is professor in the Department of Physics at North Carolina State University. He is the primary author of MVH3 (*neé* VH-1), one of the most widely used codes for astrophysical dynamics. He can be reached at blondin-at-ncsu-dot-edu.

Nathan Wichmann is currently a Senior Applications Engineer in the Cray Supercomputing Center of Excellence working with Oak Ridge National Laboratories. During his nearly ten years at Cray Inc. he has work on a range of applications include automotive, climate and weather, and fusion research. Nathan has a particular interest in compilers and processor micro-architecture and has worked in determining the direction of compiler and processor development. He can be reached at wichmann@cray.com.

References

Arnett, W. D., Bahcall, J. N., Kirshner, R. P., & Woosley, S. E. 1987, araa, 27, 701

Blondin, J. M., & Mezzacappa, A. 2007, nat, 445, 58

Bruenn, S. W. 1985, ApJS, 58, 771

Bruenn, S. W., De Nisco, K. R., & Mezzacappa, A. 2001, ApJ, 560, 326

Buras, R., Rampp, M., Janka, H.-T., & Kifonidis, K. 2006, aa, 447, 1049

Collela, P., & Woodward, P. R. 1984, J. Comp. Phys., 54, 174

Galama, T. J., Vreeswijk, P. M., van Paradijs, J., Kouveliotou, C., Augusteijn, T., Bohnhardt, H., Brewer, J. P., Doublier, V., Gonzalez, J.-F., Leibundgut, B., Lidman, C., Hainaut, O. R., Patat, F., Heise, J., in 't Zand, J., Hurley, K., Groot, P. J., Strom, R. G., Mazzali, P. A., Iwamoto, K., Nomoto, K., Umeda, H., Nakamura, T., Young, T. R., Suzuki, T., Shigeyama, T., Koshut, T., Kippen, M., Robinson, C., de Wildt, P., Wijers, R. A. M. J., Tanvir, N., Greiner, J., Pian, E., Palazzi, E., Frontera, F., Masetti, N., Nicastro, L., Feroci, M., Costa, E., Piro, L., Peterson, B. A., Tinney, C., Boyle, B., Cannon, R., Stathakis, R., Sadler, E., Begam, M. C., & Ianna, P. 1998, Nature, 395, 670

Gehrels, N., Sarazin, C. L., O'Brien, P. T., Zhang, B., Barbier, L., Barthelmy, S. D., Blustin, A., Burrows, D. N., Cannizzo, J., Cummings, J. R., Goad, M., Holland, S. T., Hurkett, C. P., Kennea, J. A., Levan, A., Markwardt, C. B., Mason, K. O., Meszaros, P., Page, M., Palmer, D. M., Rol, E.,

Sakamoto, T., Willingale, R., Angelini, L., Beardmore, A., Boyd, P. T., Breeveld, A., Campana, S., Chester, M. M., Chincarini, G., Cominsky, L. R., Cusumano, G., de Pasquale, M., Fenimore, E. E., Giommi, P., Gronwall, C., Grupe, D., Hill, J. E., Hinshaw, D., Hjorth, J., Hullinger, D., Hurley, K. C., Klose, S., Kobayashi, S., Kouveliotou, C., Krimm, H. A., Mangano, V., Marshall, F. E., McGowan, K., Moretti, A., Mushotzky, R. F., Nakazawa, K., Norris, J. P., Nousek, J. A., Osborne, J. P., Page, K., Parsons, A. M., Patel, S., Perri, M., Poole, T., Romano, P., Roming, P. W. A., Rosen, S., Sato, G., Schady, P., Smale, A. P., Sollerman, J., Starling, R., Still, M., Suzuki, M., Tagliaferri, G., Takahashi, T., Tashiro, M., Tueller, J., Wells, A. A., White, N. E., & Wijers, R. A. M. J. 2005, Nature, 437, 851

Hanhart, C., Phillips, D. R., & Reddy, S. 2001, Phys. Lett. B, 499, 9

Hannestad, S., & Raffelt, G. 1998, ApJ, 507, 339

Hix, W. R., & Thielemann, F.-K. 1999, ApJ, 511, 862

Kifonidis, K., Plewa, T., Janka, H.-T., & Müller, E. 2003, A&A, 408, 621

Lattimer, J., & Swesty, F. D. 1991, Nucl. Phys. A, 535, 331

Liebendörfer, M., Messer, O. E. B., Mezzacappa, A., Bruenn, S. W., Cardall, C. Y., & Thielemann, F.-K. 2004, ApJS, 150, 263

Marek, A., Dimmelmeier, H., Janka, H.-T., Müller, E., & Buras, R. 2006, aa, 445, 273

Matheson, T., Garnavich, P. M., Stanek, K. Z., Bersier, D., Holland, S. T., Krisciunas, K., Caldwell, N., Berlind, P., Bloom, J. S., Bolte, M., Bonanos, A. Z., Brown, M. J. I., Brown, W. R., Calkins, M. L., Challis, P., Chornock, R., Echevarria, L., Eisenstein, D. J., Everett, M. E., Filippenko, A. V., Flint, K., Foley, R. J., Freedman, D. L., Hamuy, M., Harding, P., Hathi, N. P., Hicken, M., Hoopes, C., Impey, C., Jannuzi, B. T., Jansen, R. A., Jha, S., Kaluzny, J., Kannappan, S., Kirshner, R. P., Latham, D. W., Lee, J. C., Leonard, D. C., Li, W., Luhman, K. L., Martini, P., Mathis, H., Maza, J., Megeath, S. T., Miller, L. R., Minniti, D., Olszewski, E. W., Papenkova, M., Phillips, M. M., Pindor, B., Sasselov, D. D., Schild, R., Schweiker, H., Spahr, T., Thomas-Osip, J., Thompson, I., Weisz, D., Windhorst, R., & Zaritsky, D. 2003, ApJ, 599, 394

McCray, R. 1993, ARA&A, 31, 175

Müller, E., & Steinmetz, M. 1995, Comput. Phys. Commun., 89, 45

Rauscher, T., Thielemann, F.-K., & Kratz, K.-L. 1996, MmRAS, 67, 851

Reddy, S., Prakash, M., & Lattimer, J. M. 1998, Phys. Rev. D, 58, 013009