3-D Radiative Transfer in the Next Decade

E. Baron,[1,2] Peter H. Hauschildt[3]

Abstract

Type Ia supernovae have played a crucial role in the discovery of the dark energy, via the measurement of their light curves and the determination of the peak brightness via fitting templates to the observed lightcurve shape. However, understanding the systematics of using Type Ia supernovae as cosmological probes will require 3-D radiative transfer calculations. Type II supernovae are also useful as cosmological probes, but those methods also involve radiative transfer calculations. Additionally, the connection between stripped-envelope supernovae and gamma-ray bursts strongly argue for the necessity of 3-D radiative transfer calculations. Thus, in the next decade the astrophysics community will require training in what was once the arcane field of radiative transfer as well support for the necessary hardware and training in high performance computing.

Subject headings: cosmology: dark energy — stars: atmospheres — supernovae

1. Introduction

In the last three decades or so supernova astrophysics has moved from a rather small subfield of aficionados to a major field of astrophysics. This was sparked by the discoveries of SN 1987A in 1987 and the discovery of the dark energy in 1998. Additionally, the ability to calculate 3-D hydrodynamics simulations (without true radiative transfer) has grown in the last decade. Much of what we learn in astronomy comes from spectroscopic observations, and thus in order to connect observed spectra to hydrodynamics simulations full 3-D synthetic spectra calculations are required.

Over the last 5–10 years the trend in building high performance scientific computing facilities has been to construct massively parallel systems that consist of thousands of nodes, often with several (2–8) cores per node. We refer to an individual CPU, whether it is a single

[1]Department of Physics and Astronomy, University of Oklahoma, 440 West Brooks, Rm. 100, Norman, OK 73019-2061, USA; **baron@ou.edu**

[2]Computational Research Division, Lawrence Berkeley National Laboratory, MS 50F-1650, 1 Cyclotron Rd, Berkeley, CA 94720 USA

[3]Hamburger Sternwarte, Gojenbergsweg 112, 21029 Hamburg, Germany; **yeti@hs.uni-hamburg.de**

core or one core in a multiple core node as a processing element (PE). Due to the high cost of RAM and cache, the amount of RAM per PE has remained typically in the range 256 MB – 2 GB (BlueGene, Franklin).

High performance scientific computing is driven by the need for raw speed, but also by the need for large amounts of RAM in order to attack scientific calculations in 3 dimensions with adequate resolution. For many problems the underlying physical problem is inherently local and thus large 3-D systems can be broken up into smaller problems and the memory distributed across the PEs. This is the case for pure hydrodynamics, computational fluid dynamics (CFD) and N-body calculations.

2. Radiative Transfer

Radiative Transfer is important in many areas of astrophysics: supernovae, stellar modeling, iradiated stars and planets, gamma-ray bursts, and AGN. Often one desires to do full radiation hydrodynamics but since pure 3-D hydro taxes the largest supercomputers available, coupling hydro and radiative transfer in 3-D is beyond the capabilities of current computer hardware. Radiative transfer is an inherently global problem since different physical regions are coupled by the transfer of radiation between them and the solution of the transfer equation requires a self-consistent determination of physically separate but radiatively coupled regions. Additionally, the radiative transfer equation must be solved not just in physical space, but on the contrary in the full phase space and thus the time dependent problem becomes a 7-D problem (3 spatial dimensions, 2 solid angle directional dimensions, energy or wavelength, and time; see for example Mihalas & Mihalas 1984; Mihalas 1978). For reasonable resolution, the memory requirements quickly become very large. We focus here on characteristic methods for the solution of the transfer equation, but many of our considerations are quite general, and thus also apply to, for example, Monte-Carlo and other finite difference methods as well.

In the last twenty years, numerical methods have been developed to treat the fully non-LTE (NLTE) radiative transfer equation in 1-D relativistic flows using very large model atoms and extremely large databases for molecules. Modern numerical methods allow for efficient parallelization.

Due to the global nature of the RT problem each processor that calculates radiative transfer needs to keep information on the entire physical grid. If a domain decomposion of the spatial grid is used to spread the data between different PE, different solid angles require either massive communication between different PEs or domain adjustments during the formal solution. This is theoretically possible, but for straightforward implementation the amount of communication is prohibitive.

The radiative transfer problem is solved when the mean intensity, J, is known. Thus, hydro codes which claim to include radiative transfer, but specify the value of J *a priori* do not solve the radiative transfer problem. The mean intensity, J, is obtained from the source function S by a formal solution of the RTE which is symbolically written using the Λ-operator Λ as

$$J = \Lambda S. \tag{1}$$

The source function is given by $S = (1-\epsilon)J + \epsilon B$, where ϵ denotes the thermal coupling parameter and B is Planck's function.

The Λ-iteration method, i.e. to solve Eq. 1 by a fixed-point iteration scheme of the form

$$J_{\rm new} = \Lambda S_{\rm old}, \quad S_{\rm new} = (1-\epsilon)J_{\rm new} + \epsilon B,$$

fails in the case of large optical depths and small ϵ.

The idea of the ALI or operator splitting (OS) method (Olson et al. 1986; Olson & Kunasz 1987; Scharmer 1984) is to reduce the eigenvalues of the amplification matrix in the iteration scheme by introducing an approximate Λ-operator (ALO) Λ^* and to split Λ according to

$$\Lambda = \Lambda^* + (\Lambda - \Lambda^*)$$

and rewrite Eq. 2 as

$$J_{\rm new} = \Lambda^* S_{\rm new} + (\Lambda - \Lambda^*)S_{\rm old}.$$

To achieve a significant improvement compared to the Λ-iteration, the operator Λ^* is constructed so that the eigenvalues of the iteration matrix are much smaller than unity, resulting in swift convergence. Using parts of the exact Λ matrix (e.g., its diagonal or a tri-diagonal form) will optimally reduce the eigenvalues of the iteration matrix. Non-local Λ^* (non-diagonal) operators lead to excellent convergence rates and avoid the problem of false convergence that is inherent in the Λ iteration method and can also be an issue for diagonal (purely local) Λ^* operators.

The largest storage in the 3-D RT is needed for the Λ^* operator, which propagates the global solution. In our implementation this operator consists of diagonal terms and terms from nearest neighbors which at each physical point (voxel) will be $3^3 = 27$ eight-byte numbers. Additional storage should be less than a total of 64 eight-byte numbers. The additional storage cost is repaid many times over with the acceleration of the rate of convergence, dramatically reducing the computation time. Typical spatial grids will be of order $128^3 \sim 10^6$ points — 10^9 points could easily be needed, but that is far beyond the reach of hardware that could be available in the next 10 years (see Table 1). The need to store the entire grid in memory is due to the nature of the problem and not just the numerical algorithm we are discussing, 3-D Monte-Carlo schemes have the exact same restriction.

Just for the radiative transfer, the storage requirements are not too severe; however, the more severe storage requirement is that the values of the number densities of each species must be stored in order to calculate the wavelength dependent opacities at each wavelength point. In a typical calculation the total number of species (ions and molecules) will be of the order of one thousand. For the radiative transfer only total opacity at each wavelength is required and hence one could recompute the number density at each wavelength point. However, typical applications require from 10,000 to 1,000,000 wavelength points and thus this memory savings would come at an enormous cost of CPU time as a solution of the EOS is costly. This storage/computational tradeoff may be mitigated by a clever table interpolation scheme. NLTE lines and the rate operator expand the storage per point to between 3–30 MB. Multiplying these numbers we come to enormous storage requirements, but these can be made tractable by domain decomposition for the structural data (EOS, opacities, NLTE). These unsolved problems point to fruitful collaboration between astronomers and computer scientists.

3. Applications

Table 1 shows the number of PEs needed to calculate standard astrophysical objects. It is clear that just to handle the physical data large numbers of PEs are required. Further parallelization over solid angle is easily implemented (see below) and can increase the number of PEs required significantly in order to keep wall-clock times below one day.

Figure 1 show visualizations of the results for 3D continuum and line transfer. The importance of scattering is specified by the "thermalization parameter" $\epsilon = \kappa/(\kappa + \sigma)$ where κ is the absorption opacity and σ is the scattering opacity. The RT problem was solved for $\epsilon = 1$ (left panel) and 10^{-4} (right panel) and a formal solution with the converged source functions was computed for given viewing angles. The graphs are actual images of the intensities as they would be seen by an external observer different angles. The visible surface is to the left, the 'sides' of the computational box could not be seen by an observer and are shown for information only. The effect of scattering on the images is similar to terrestrial fog in that it reduces the contrast of visible features; even moderate scattering of $\epsilon_c = 10^{-4}$ significantly reduces visibility. Limb darkening is also clearly visible in the figures.

Figure 2 shows the wall-clock time as a function of resolution (number of momentum-space angles n_θ and n_ϕ or number of CPUs). The computational work per CPU is kept constant with each CPU required to calculate 16 characteristics. The modest (14%) increase in wall-clock time from 16 CPUs to 16,384 CPUs is acceptable given the huge increase in communication required and the fact that load balancing is quite simple.

4. Conclusions

The advent of 3-D NLTE radiative transfer will significantly enhance our ability to analyze and interpret both observed spectra and theoretical hydrodynamics calculations. 3-D radiative transfer has already shown its potential for new results in the current reconsideration of solar abundances. Detailed 3-D radiative transfer will shed further light on this problem soon, as well as on problems in iradiated stars and planets, supernovae, gamma-ray bursts, AGN, and cosmology. Thus, both man-power and computational resources will be required for this reinvigorated field in the next decade.

REFERENCES

Mihalas, D. 1978, Stellar Atmospheres (New York: W. H. Freeman)

Mihalas, D. & Mihalas, B. W. 1984, Foundations of Radiation Hydrodynamics (Oxford: Oxford University)

Olson, G. L., Auer, L. H., & Buchler, J. R. 1986, JQSRT, 35, 431

Olson, G. L. & Kunasz, P. B. 1987, JQSRT, 38, 325

Scharmer, G. B. 1984, in Methods in Radiative Transfer, ed. W. Kalkofen (Cambridge: Cambridge Univ. Press), 173

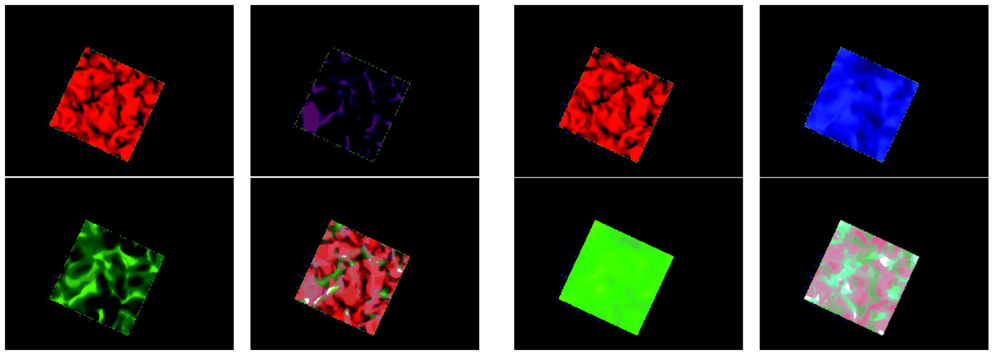

Fig. 1.— Visualization of the results for the line 3D radiation transfer with $\epsilon_l = 1$ (left) and 10^{-4} (right). The images are intensities in the directions $\phi = 25\,\mathrm{deg}$ and $\theta = 0\,\mathrm{deg}$. The top left panel is the image in the continuum, the top right panel the image at the line center, the bottom left panel the image in the line wing, the bottom right panel is a composite image.

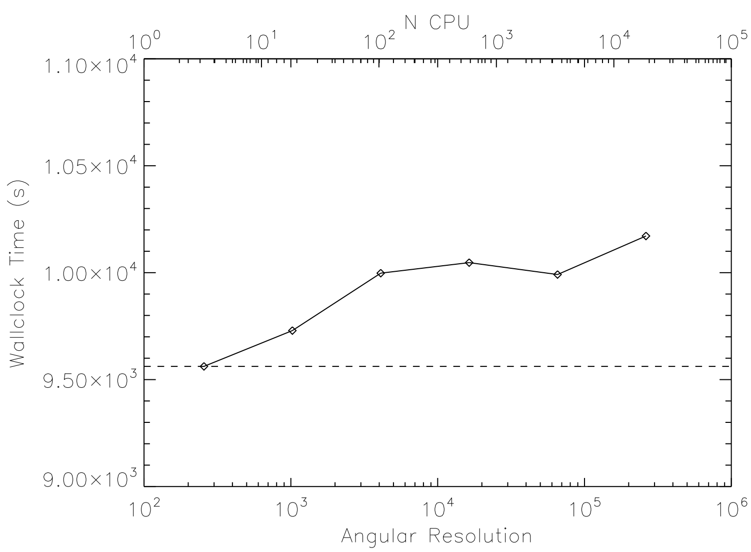

Fig. 2.— The wall-clock time to solve a scattering line $\epsilon = 0.1$ on the the NERSC Franklin Cray-XT4 as a function of momentum frame angular resolution. The test was run so the amount of computational work per processor was constant. The dashed horizontal line indicates perfect scaling. The roughly 14% communication time increase from 16 to 16,384 processors is acceptable.

Model	2*nx+1	2*ny+1	2*nz+1	N_{points}	RT [MB]	EOS [MB]	NLTE [MB]	RAM/PR [MB]	PE/RT	Free/Core [MB]	PEs Reqd
Sun (1D)	64	1	1	64	0	0	376	2048	2	2048	1
Sun	128	128	128	2×10^9	1024	16000	12×10^9	2048	2	1536	8032
Supernova	1000	32	64	2×10^9	1000	15625	12×10^9	2048	2	1548	7783

Table 1: The number of PEs required in order to calculate 3-D NLTE problems. For the Sun we assume that only a small fraction will be calculated using periodic boundary conditions as shown in Figure 1. For supernovae we assume that a spherical coordinate system will be used so that the "X-direction" corresponds to radial zones and the Y and Z directions correspond to coordinate angles.