

Porting and Optimizing VERTEX-PROMETHEUS on the Cray XE6 at HLRS for Three-Dimensional Simulations of Core-Collapse Supernova Explosions of Massive Stars

Florian Hanke, Bernhard Müller, Hans-Thomas Janka
*Max-Planck-Institut für Astrophysik
Karl-Schwarzschild-Strasse 1
85748 Garching bei München, Germany
{fhanke,bjmuellr;thj}@mpa-garching.mpg.de*

Andreas Marek
*Rechenzentrum Garching
Boltzmannstrasse 2
85748 Garching bei München, Germany
amarek@rzg.mpg.de*

Abstract

Supernova explosions are among the most powerful cosmic events, whose physical mechanism and consequences are still incompletely understood. We have developed a fully MPI-OpenMP parallelized version of our VERTEX-PROMETHEUS code in order to perform three-dimensional simulations of stellar core-collapse and explosion on Tier-0 systems such as Hermit at HLRS. Tests on up to 64000 cores have shown excellent scaling behavior. In this paper we present the system of equations and the algorithm for its solution that are employed in our code VERTEX-PROMETHEUS. We also discuss the parallelization of VERTEX-PROMETHEUS and present our progress in porting, optimizing, and performing production runs on a large variety of machines, starting from vector machines and reaching to modern systems. In particular the results of our efforts to achieve good parallel scaling on the new Cray XE6 at HLRS Stuttgart are highlighted.

Keywords

Supernova explosions, code porting, optimizing,

I. INTRODUCTION

A star more massive than about 8 solar masses ends its life in a catastrophic explosion, a supernova. Its quiescent evolution comes to an end, when the pressure in its inner layers is no longer able to balance the inward pull of gravity. Throughout its life, the star sustained this balance by generating energy through a sequence of nuclear fusion reactions, forming increasingly heavier elements in its core. However, when the core consists mainly of iron-group nuclei, central energy generation ceases. The fusion reactions producing iron-group nuclei relocate to the core's surface, and their "ashes" continuously increase the core's mass. Similar to a white dwarf, such a core is stabilised against gravity by the pressure of its degenerate gas of electrons. However, to remain stable, its mass must stay smaller than (roughly) the Chandrasekhar limit. When the core grows larger than this limit, it collapses to a neutron star, and a huge amount ($\sim 10^{53}$ erg) of gravitational binding energy is set free. Most ($\sim 99\%$) of this energy is radiated away in neutrinos, but a small fraction is transferred to the outer stellar layers and drives the violent mass ejection, which disrupts the star in a supernova.

Despite 40 years of research, the details of how this energy transfer happens and how the explosion is initiated are still not well understood. Observational evidence about the physical processes deep inside the collapsing star is sparse and almost exclusively indirect. The only direct observational access is via measurements of neutrinos or gravitational waves. To obtain insight into the events in the core, one must therefore heavily rely on sophisticated numerical simulations. The enormous amount of computer power required for this purpose has led to the use of several, often questionable, approximations and numerous ambiguous results in the past. Fortunately, however, the development of numerical tools and computational resources has meanwhile advanced to a point, where it is becoming possible to perform multi-dimensional simulations with unprecedented accuracy. Therefore there is hope that the physical processes which are essential for the explosion can finally be unravelled.

An understanding of the explosion mechanism is required to answer many important questions of nuclear, gravitational, and astro-physics like the following:

- How do the explosion energy, the explosion timescale, and the mass of the compact remnant depend on the progenitor's mass? Is the explosion mechanism the same for all progenitors? For which stars are black holes left behind as compact remnants instead of neutron stars?
- What is the role of the – incompletely known – equation of state (EoS) of the proto-neutron star? Do softer or stiffer EoSs favour the explosion of a core collapse supernova?
- How do neutron stars receive their natal kicks? Are they accelerated by asymmetric mass ejection and/or anisotropic neutrino emission?
- What are the generic properties of the neutrino emission and of the gravitational wave signal that are produced during stellar core collapse and explosion? Up to which distances could these signals be measured with operating

or planned detectors on earth and in space? And what can one learn about supernova dynamics or nuclear and particle physics from a future measurement of such signals in the case of a Galactic supernova?

- How do supernovae contribute to the enrichment of the intergalactic medium with heavy elements? What kind of nucleosynthesis processes occur during and after the explosion? Can the elemental composition of supernova remnants be explained correctly by the numerical simulations? Does the rapid neutron capture process (r-process), which produces e.g. gold and the actinides, take place in supernovae?

II. NUMERICAL MODELING

A. History and constraints

According to theory, a shock wave is launched at the moment of “core bounce” when the neutron star begins to emerge from the collapsing stellar iron core. There is general agreement, supported by all “modern” numerical simulations, that this shock is unable to propagate directly into the stellar mantle and envelope, because it loses too much energy in dissociating iron into free nucleons while it moves through the outer core. The “prompt” shock ultimately stalls. Thus the currently favoured theoretical paradigm exploits the fact that a huge energy reservoir is present in the form of neutrinos, which are abundantly emitted from the hot, nascent neutron star. The absorption of electron neutrinos and anti-neutrinos by free nucleons in the post-shock layer is thought to reenergize the shock, thus triggering the supernova explosion.

Detailed *spherically symmetric* hydrodynamic models, which recently include a very accurate treatment of the time-dependent, multi-flavour, multi-frequency neutrino transport based on a numerical solution of the Boltzmann transport equation [1], [2], reveal that this “delayed, neutrino-driven mechanism” does not work as simply as originally envisioned. Although in principle able to trigger the explosion (e.g., [3], [4], [5]), neutrino energy transfer to the post-shock matter turned out to be too weak. For inverting the infall of the stellar core and initiating powerful mass ejection, an increase of the efficiency of neutrino energy deposition is needed.

A number of physical phenomena have been pointed out that can enhance neutrino energy deposition behind the stalled supernova shock. They are all linked to the fact that the real world is multi-dimensional instead of spherically symmetric (or one-dimensional; 1D) as assumed in the works cited above:

- (1) Convective instabilities in the neutrino-heated layer between the neutron star and the supernova shock develop to violent convective overturn [6]. This convective overturn is helpful for the explosion, mainly because (a) neutrino-heated matter rises and increases the pressure behind the shock, thus pushing the shock further out, (b) cool matter is able to penetrate closer to the neutron star where it can absorb neutrino energy more efficiently, and (c) the rise of freshly heated matter reduces energy losses by the reemission of neutrinos. These effects allow multi-dimensional models to explode easier than spherically symmetric ones [7], [8], [9].
- (2) Recent work [10], [11], [12], [13] has demonstrated that the stalled supernova shock is also subject to a second non-radial low-mode instability, called the standing accretion shock instability or “SASI” for short, which can grow to a dipolar, global deformation of the shock [12], [14], [15].
- (3) Convective energy transport inside the nascent neutron star [16], [17], [18] might enhance the energy transport to the neutrinosphere and could thus boost the neutrino luminosities. This would in turn increase the neutrino-heating behind the shock.

This list of multi-dimensional phenomena (limited to non-magnetized supernova cores) awaits more detailed exploration by multi-dimensional simulations. Until recently, such simulations have been performed with only a grossly simplified treatment of the involved microphysics, in particular of the neutrino transport and neutrino-matter interactions. At best, grey (i.e., single energy) flux-limited diffusion schemes were employed. Since, however, the role of the neutrinos is crucial for the problem, and because previous experience shows that the outcome of simulations is indeed very sensitive to the employed transport approximations, studies of the explosion mechanism require the best available description of the neutrino physics. This implies that one has to solve the Boltzmann transport equation for neutrinos.

B. The mathematical model

As core-collapse supernovae involve such a complex interplay of hydrodynamics, self-gravity and neutrino heating and cooling, numerical modellers face a classical “multiphysics” problem. Although the overall problem can still be formulated as a system of non-linear partial differential equations, rather dissimilar methods – sometimes with conflicting requirements on the computer architecture and the parallelization strategy – need to be applied to treat individual subsystems. In the case of our code, the system of equations that needs to be solved consists of the following components:

- The multi-dimensional Euler equations of (relativistic) hydrodynamics, supplemented by advection equations for the electron fraction and the chemical composition of the fluid, and formulated in spherical polar coordinates;
- equations for the space-time metric (or in the Newtonian case, the Poisson equation) for calculating the gravitational source terms in the Euler equations;
- the Boltzmann transport equation and/or its moment equations which determine the (non-equilibrium) distribution function of the neutrinos;

- the emission, absorption, and scattering rates of neutrinos, which are required for the solution of the neutrino transport equations;
- the equation of state of the stellar fluid, which provides the closure relation between the variables entering the Euler equations, i.e. density, momentum, energy, electron fraction, composition, and pressure.

In what follows we will briefly summarise the neutrino transport algorithms, thus focusing on the major computational kernel of our code. For a more complete description of the entire code we refer the reader to [19], [20], and the references therein.

C. “Ray-by-ray plus” method for the neutrino transport problem

The crucial quantity required to determine the source terms for the energy, momentum, and electron fraction of the fluid owing to its interaction with the neutrinos is the neutrino distribution function in phase space, $f(r, \vartheta, \phi, \epsilon, \Theta, \Phi, t)$. Equivalently, the neutrino intensity $I = c/(2\pi\hbar c)^3 \cdot \epsilon^3 f$ may be used. Both are time-dependent functions in a six-dimensional phase space, as they describe, at every point in space (r, ϑ, ϕ) , the distribution of neutrinos propagating with energy ϵ into the direction (Θ, Φ) at time t (Fig. 1).

The evolution of I (or f) in time is governed by the Boltzmann equation, and solving this equation is, in general, a six-dimensional problem (as time is usually not counted as a separate dimension). A solution of this equation by direct discretization (using an S_N scheme) would require computational resources in the PetaFlop range. Although there are attempts by at least one group in the United States to follow such an approach, we feel that, with the currently available computational resources, it is mandatory to reduce the dimensionality of the problem.

Actually this should be possible, since the source terms entering the hydrodynamic equations are *integrals* of I over momentum space (i.e. over ϵ , Θ , and Φ), and thus only a fraction of the information contained in I is truly required to compute the neutrino effects on the dynamics of the flow. It therefore makes sense to consider angular moments of I , and to solve evolution equations for these moments, instead of dealing with the Boltzmann equation directly. The 0th to 3rd order moments are defined as

$$\mathbf{J}, \mathbf{H}, \mathbf{K}, \mathbf{L}, \dots(r, \vartheta, \phi, \epsilon, t) = \frac{1}{4\pi} \int I(r, \vartheta, \phi, \epsilon, \Theta, \Phi, t) \mathbf{n}^{0,1,2,3,\dots} d\Omega \quad (1)$$

where $d\Omega = \sin \Theta d\Theta d\Phi$, $\mathbf{n} = (\cos \Theta, \sin \Theta \cos \Phi, \sin \Theta \sin \Phi)$, and exponentiation represents repeated application of the dyadic product. Note that the moments are *tensors* of the required rank.

So far no approximations have been made. In order to reduce the size of the problem even further, one needs to resort to assumptions on its symmetry. At this point, one assumes that I is independent of Θ and Φ , then each of the angular moments of I becomes a *scalar*, which depends on three spatial dimensions, and one dimension in momentum space: $J, H, K, L = J, H, K, L(r, \vartheta, \phi, \epsilon, t)$. Thus the neutrino moment equations at different angular directions (except for some terms which can be accounted for explicitly in an operator split) decouple from each other. Therefore, for each “radial ray”, i.e. for all zones of same angle, the moment equations can be solved independently. Except for some additional terms, this problem is identical to solving $N_\theta \times N_\phi$ times the moment equations for a spherically symmetric star with $N_\theta \times N_\phi$ being the number of grid zones in polar direction. As we will explain later, the great advantage of our “ray-by-ray” neutrino transport is the easy way to obtain perfect scaling behaviour to a large number of cores.

1) *The system of equations:* With the aforementioned assumptions it can be shown [19], that in the Newtonian approximation the following two transport equations need to be solved in order to compute the source terms for the




Figure 1. Illustration of the phase space coordinates (see the main text).

energy and electron fraction of the fluid:

$$\begin{aligned}
& \left(\frac{1}{c} \frac{\partial}{\partial t} + \beta_r \frac{\partial}{\partial r} + \frac{\beta_\vartheta}{r} \frac{\partial}{\partial \vartheta} + \frac{\beta_\varphi}{r \sin \vartheta} \frac{\partial}{\partial \varphi} \right) J + J \left(\frac{1}{r^2} \frac{\partial(r^2 \beta_r)}{\partial r} + \frac{1}{r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} + \frac{1}{r \sin \vartheta} \frac{\partial \beta_\varphi}{\partial \varphi} \right) \\
& + \frac{1}{r^2} \frac{\partial(r^2 H)}{\partial r} + \frac{\beta_r}{c} \frac{\partial H}{\partial t} - \frac{\partial}{\partial \epsilon} \left\{ \frac{\epsilon}{c} \frac{\partial \beta_r}{\partial t} H \right\} - \frac{\partial}{\partial \epsilon} \left\{ \epsilon J \left(\frac{\beta_r}{r} + \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} + \frac{1}{2r \sin \vartheta} \frac{\partial \beta_\varphi}{\partial \varphi} \right) \right\} \\
& - \frac{\partial}{\partial \epsilon} \left\{ \epsilon K \left(\frac{\partial \beta_r}{\partial r} - \frac{\beta_r}{r} - \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} - \frac{1}{2r \sin \vartheta} \frac{\partial \beta_\varphi}{\partial \varphi} \right) \right\} + J \left(\frac{\beta_r}{r} + \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} + \frac{1}{2r \sin \vartheta} \frac{\partial \beta_\varphi}{\partial \varphi} \right) \\
& + K \left(\frac{\partial \beta_r}{\partial r} - \frac{\beta_r}{r} - \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} - \frac{1}{2r \sin \vartheta} \frac{\partial \beta_\varphi}{\partial \varphi} \right) + \frac{2}{c} \frac{\partial \beta_r}{\partial t} H = C^{(0)}, \quad (2)
\end{aligned}$$

$$\begin{aligned}
& \left(\frac{1}{c} \frac{\partial}{\partial t} + \beta_r \frac{\partial}{\partial r} + \frac{\beta_\vartheta}{r} \frac{\partial}{\partial \vartheta} + \frac{\beta_\varphi}{r} \frac{\partial}{\partial \varphi} \right) H + H \left(\frac{1}{r^2} \frac{\partial(r^2 \beta_r)}{\partial r} + \frac{1}{r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} + \frac{1}{r \sin \vartheta} \frac{\partial \beta_\varphi}{\partial \varphi} \right) \\
& + \frac{\partial K}{\partial r} + \frac{3K - J}{r} + H \left(\frac{\partial \beta_r}{\partial r} \right) + \frac{\beta_r}{c} \frac{\partial K}{\partial t} - \frac{\partial}{\partial \epsilon} \left\{ \frac{\epsilon}{c} \frac{\partial \beta_r}{\partial t} K \right\} \\
& - \frac{\partial}{\partial \epsilon} \left\{ \epsilon L \left(\frac{\partial \beta_r}{\partial r} - \frac{\beta_r}{r} - \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} - \frac{1}{2r \sin \vartheta} \frac{\partial \beta_\varphi}{\partial \varphi} \right) \right\} \\
& - \frac{\partial}{\partial \epsilon} \left\{ \epsilon H \left(\frac{\beta_r}{r} + \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} + \frac{1}{2r \sin \vartheta} \frac{\partial \beta_\varphi}{\partial \varphi} \right) \right\} + \frac{1}{c} \frac{\partial \beta_r}{\partial t} (J + K) = C^{(1)}. \quad (3)
\end{aligned}$$

These are evolution equations for the neutrino energy density, J , and the neutrino flux, H , and follow from the zeroth and first moment equations of the comoving frame (Boltzmann) transport equation in the Newtonian, $\mathcal{O}(v/c)$ approximation. The quantities $C^{(0)}$ and $C^{(1)}$ are source terms that result from the collision term of the Boltzmann equation, while $\beta_r = v_r/c$, $\beta_\vartheta = v_\vartheta/c$, and $\beta_\varphi = v_\varphi/c$, where v_r , v_ϑ , and v_φ are the components of the hydrodynamic velocity, and c is the speed of light. The functional dependencies $\beta_r = \beta_r(r, \vartheta, \varphi, t)$, $J = J(r, \vartheta, \varphi, \epsilon, t)$, etc. are suppressed in the notation. This system includes four unknown moments (J, H, K, L) but only two equations, and thus needs to be supplemented by two more relations. This is done by substituting $K = f_K \cdot J$ and $L = f_L \cdot J$, where f_K and f_L are the variable Eddington factors, which for the moment may be regarded as being known, but in our case are indeed determined from a separate simplified (“model”) Boltzmann equation.

The moment equations (2) and (3) are very similar to the $\mathcal{O}(v/c)$ equations in spherical symmetry which were solved in the 1D simulations of [21] (see Eqs. (7),(8),(30), and (31) of the latter work). This similarity has allowed us to reuse a good fraction of the one-dimensional version of the transport part, for coding the multi-dimensional algorithm. The additional terms necessary for this purpose have been set in boldface above.

Finally, the changes of the energy, e , and electron fraction, Y_e , required for the hydrodynamics are given by the following two equations

$$\frac{de}{dt} = -\frac{4\pi}{\rho} \int_0^\infty d\epsilon \sum_{\nu \in (\nu_e, \bar{\nu}_e, \dots)} C_\nu^{(0)}(\epsilon), \quad (4)$$

$$\frac{dY_e}{dt} = -\frac{4\pi m_B}{\rho} \int_0^\infty \frac{d\epsilon}{\epsilon} \left(C_{\nu_e}^{(0)}(\epsilon) - C_{\bar{\nu}_e}^{(0)}(\epsilon) \right) \quad (5)$$

(for the momentum source terms due to neutrinos see [19]). Here m_B is the baryon mass, and the sum in Eq. (4) runs over all neutrino types. The full system consisting of Eqs. (2–5) is stiff, and thus requires an appropriate discretization scheme for its stable solution.

2) Method of solution: In order to discretize Eqs. (2–5), the spatial domain $[0, r_{\max}] \times [\vartheta_{\min}, \vartheta_{\max}] \times [\varphi_{\min}, \varphi_{\max}]$ is covered by N_r radial, N_ϑ latitudinal, and N_φ longitudinal zones, where $\vartheta_{\min} = 0$ and $\vartheta_{\max} = \pi$ correspond to the north and south poles, respectively, of the spherical grid and $\varphi_{\min} = 0$ and $\varphi_{\max} = 2\pi$ covers the full sphere. (In general, we allow for grids with different radial resolutions in the neutrino transport and hydrodynamic parts of the code. The number of radial zones for the hydrodynamics will be denoted by N_r^{hyd} .) The number of bins used in energy space is N_ϵ and the number of neutrino types taken into account is N_ν .

The equations are solved in three operator-split steps corresponding to a lateral, an azimuthal and a radial sweep.

In the first two steps, we treat the boldface terms in the respectively first lines of Eqs. (2–3), which describe the lateral and azimuthal advection of the neutrinos with the stellar fluid, and thus couple the angular moments of the neutrino distribution of neighbouring angular zones. For this purpose we consider the equations

$$\frac{1}{c} \frac{\partial \Xi}{\partial t} + \frac{1}{r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta \Xi)}{\partial \vartheta} = 0, \quad (6)$$

$$\frac{1}{c} \frac{\partial \Xi}{\partial t} + \frac{1}{r \sin \vartheta} \frac{\partial (\beta_\varphi \Xi)}{\partial \varphi} = 0, \quad (7)$$

where Ξ represents one of the moments J or H . Although it has been suppressed in the above notation, an equation of this form has to be solved for each radius, for each energy bin, and for each type of neutrino. An explicit upwind scheme is used for this purpose.

In the third step, the radial sweep is performed. Several points need to be noted here:

- Terms in boldface not yet taken into account in the lateral sweep, need to be included into the discretization scheme of the radial sweep. This can be done in a straightforward way since these remaining terms do not include derivatives of the transport variables J or H . They only depend on the hydrodynamic velocities v_ϑ and v_φ , which are a *constant* scalar field for the transport problem.
- The right hand sides (source terms) of the equations and the coupling in energy space have to be accounted for. The coupling in energy is non-local, since the source terms of Eqs. (2) and (3) stem from the Boltzmann equation, which is an integro-differential equation and couples all the energy bins.
- The discretization scheme for the radial sweep is *implicit* in time. Explicit schemes would require very small time steps to cope with the stiffness of the source terms in the optically thick regime, and the small CFL time step dictated by neutrino propagation with the speed of light in the optically thin regime. Still, even with an implicit scheme $\gtrsim 10^5$ time steps are required per simulation. This makes the calculations expensive.

Once the equations for the radial sweep have been discretized in radius and energy, the resulting solver is applied ray-by-ray for each pair of angles (ϑ, φ) and for each type of neutrino; i.e. for constant (ϑ, φ) , N_ν two-dimensional problems need to be solved.

The discretization itself is done using a second order accurate scheme with backward differencing in time according to [21]. This leads to a non-linear system of algebraic equations, which is solved by Newton-Raphson iteration with explicit construction and inversion of the corresponding Jacobian matrix with the Block-Thomas algorithm.




Figure 2. Schematic sketch of the MPI communication pattern for an angular direction in the hydrodynamics part of our code. The red rectangle symbolizes the data available in a specific MPI task, the surrounding area marked by the dashed line reflects the zones to be communicated.

III. PORTING AND SCALING ON THE CRAY XE6 “HERMIT” AT HLRS

A. Parallelization strategy

The ray-by-ray approximation readily lends itself to parallelization over the different angular zones. In order to make efficient use of modern supercomputer systems with relatively small shared-memory units (e.g. 16 CPUs per node on Cray XE6), distributed memory parallelism is indispensable. An MPI version of the VERTEX-PROMETHEUS code using domain decomposition was initially developed within a cooperation between MPA and the Teraflop Workbench at the HLRS in 2007/2008. Since then, the parallelization of VERTEX-PROMETHEUS has been further extended to allow good scaling on several thousands of cores as required for future 3D supernova simulations.

The VERTEX-PROMETHEUS code employs a hybrid MPI-OpenMP parallelization scheme, in which the parallelization of the transport module – the main computational kernel and most CPU-intense part of the code – is along radial “rays” for fixed angular bins of the three-dimensional grid. Hence, every “ray” of the transport is treated by one core using as many OpenMP threads as cores available on an individual node. This strategy allows almost perfect scaling behavior, since almost no MPI communication is necessary between individual rays during the transport step.

The MPI-parallelization of the much less expensive hydrodynamical part PROMETHEUS is based on standard domain decomposition methods. Hereby, the reconstruction scheme used to solve the hydrodynamic equations requires so-called “ghost-zones”, which have to be available in each MPI task. In our case, 4 ghost zones are required on each cell interface

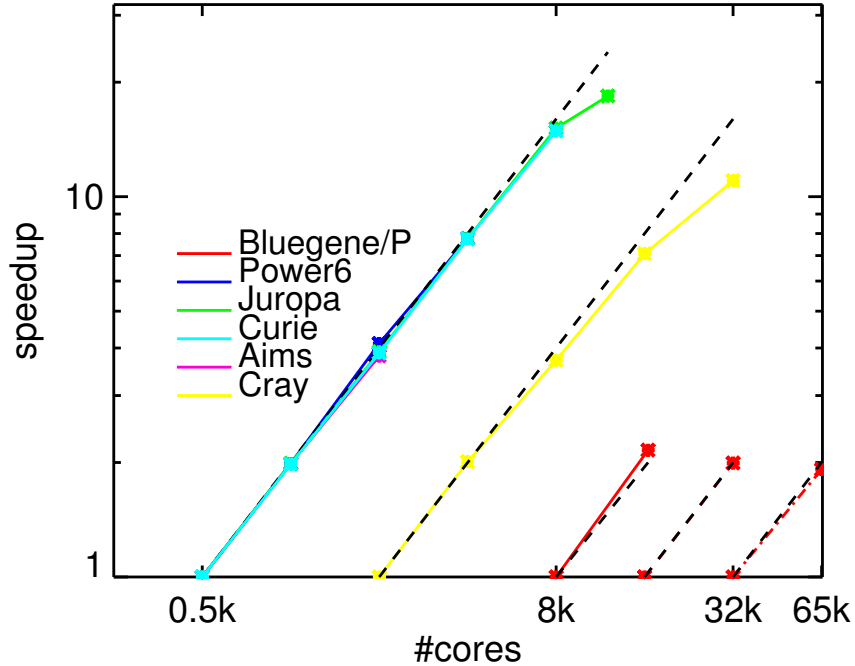


Figure 3. Strong scaling of VERTEX-PROMETHEUS on different machines and architectures. The colored lines show the speedup on the respective machine relative to the run with the smallest number of cores for a given problem size. The symbols mark the number of cores on which the timings were done. The dashed black lines indicate a theoretically perfect scaling behavior. Note that several lines lie on top of each other. Also note that due to limited memory on BlueGene/P three different setups were timed, which are shown as three separate speedup curves.

in angular directions to integrate one time step and these zones have to be MPI communicated to the neighbouring MPI tasks. A sketch of grid zones to be MPI communicated is illustrated in Figure 2.

B. Porting VERTEX-PROMETHEUS to the Cray XE6 “HERMIT”

As demonstrated in Figure 3, we have already obtained excellent scaling behavior with the explained parallelization strategy. For example, we have performed scaling tests on the BlueGene/P system JUGENE at the Forschungszentrum Jülich to demonstrate that our VERTEX-PROMETHEUS code scales perfectly up to 65,000 cores.

Since our VERTEX-PROMETHEUS code runs successfully on several architectures, the code should in principle work out of the box. However, we had to change several smaller statements in order to be able to compile the code. Furthermore, while performing the first scaling test on the Cray of the HLRS we detected that the routine, which calculates the most important neutrino interaction rates, shows poor performance. Initially, we have used the same version of this part of the transport solver, which performs perfectly using the Intel compiler. To obtain better results on the Cray XE6 we have rewritten this routine and we use now a vectorized version with one main loop.

Employing this single optimization, the code scales well on up to 32,000 cores of the Cray XE6 at HLRS as shown in Figure 3. However, the scaling behavior is still slightly worse than on Intel platforms. We plan to analyse the detailed code performance on the Cray XE6 further to get better results of the scaling tests.

Another point concerning the special characteristics of the Cray XE6 is the strong interconnection of the individual nodes. We cannot profit a lot by this feature since our code needs only a low amount of communication (less than 5% of the total computing time).

Furthermore, we want to improve the performance of I/O on the Cray XE6. The I/O is now handled by means of parallel HDF5 to ensure high scalability and to eliminate the excessive memory consumption associated with temporary I/O arrays on the root node. The handling of I/O performs quite well on IBM BlueGene and Intel systems, however we want to optimize I/O on the Cray XE6 further.

IV. CONCLUSION

We have presented our main simulation tool VERTEX-PROMETHEUS. In the past years, we have developed a fully MPI/OpenMP parallelized code version to be able to perform large scale runs on several thousand cores. At the moment our code shows excellent scaling behavior on several platforms. After the new Cray XE6 “HERMIT” had become available at HLRS, we have ported VERTEX-PROMETHEUS to this new system. With minor optimizations (required by the compiler) the code scales now up to 32,000 cores.

Since our code is now ready to run on the new Cray XE6 at HRLS, we are ready to start the first generation of three-dimensional simulations of core-collapse supernova explosions this year. This simulations are extremely expensive

(several 10^{20} floating point operations) that we need to strongly rely on Tier-0 systems such as “HERMIT”. Only systems like the new Cray XE6 in Stuttgart give us the possibility to advance our understanding of the details of the explosions mechanism of core-collapse supernovae.

ACKNOWLEDGMENT

We thank especially K. Benkert for her extremely valuable and fruitful work on the MPI version of VERTEX. Support by the Deutsche Forschungsgemeinschaft through the SFB/TR27 “Neutrinos and Beyond” and the SFB/TR7 “Gravitational Wave Astronomy”, and by the Cluster of Excellence EXC 153 “Origin and Structure of the Universe” (<http://www.universe-cluster.de>) are acknowledged, as well computer time grants of the HLRS, NIC Jülich, and Rechenzentrum Garching are acknowledged.

REFERENCES

- [1] Rampp, M., Janka, H.T.: Spherically Symmetric Simulation with Boltzmann Neutrino Transport of Core Collapse and Postbounce Evolution of a $15 M_{\odot}$ Star. *Astrophys. J.* **539** (2000) L33–L36
- [2] Liebendörfer, M., Mezzacappa, A., Thielemann, F., Messer, O.E., Hix, W.R., Bruenn, S.W.: Probing the gravitational well: No supernova explosion in spherical symmetry with general relativistic Boltzmann neutrino transport. *Phys. Rev. D* **63** (2001) 103004+
- [3] Bethe, H.A.: Supernova mechanisms. *Reviews of Modern Physics* **62** (1990) 801–866
- [4] Burrows, A., Goshy, J.: A Theory of Supernova Explosions. *Astrophys. J.* **416** (1993) L75
- [5] Janka, H.T.: Conditions for shock revival by neutrino heating in core-collapse supernovae. *Astron. Astrophys.* **368** (2001) 527–560
- [6] Herant, M., Benz, W., Colgate, S.: Postcollapse hydrodynamics of SN 1987A - Two-dimensional simulations of the early evolution. *Astrophys. J.* **395** (1992) 642–653
- [7] Herant, M., Benz, W., Hix, W.R., Fryer, C.L., Colgate, S.A.: Inside the supernova: A powerful convective engine. *Astrophys. J.* **435** (1994) 339
- [8] Burrows, A., Hayes, J., Fryxell, B.A.: On the nature of core-collapse supernova explosions. *Astrophys. J.* **450** (1995) 830
- [9] Janka, H.T., Müller, E.: Neutrino heating, convection, and the mechanism of Type-II supernova explosions. *Astron. Astrophys.* **306** (1996) 167+
- [10] Thompson, C.: Accretional Heating of Asymmetric Supernova Cores. *Astrophys. J.* **534** (2000) 915–933
- [11] Blondin, J.M., Mezzacappa, A., DeMarino, C.: Stability of Standing Accretion Shocks, with an Eye toward Core-Collapse Supernovae. *Astrophys. J.* **584** (2003) 971–980
- [12] Scheck, L., Plewa, T., Janka, H.T., Kifonidis, K., Müller, E.: Pulsar Recoil by Large-Scale Anisotropies in Supernova Explosions. *Phys. Rev. Letters* **92** (2004) 011103+
- [13] Foglizzo, T., Galletti, P., Scheck, L., Janka, H.T.: Instability of a Stalled Accretion Shock: Evidence for the Advective-Acoustic Cycle. *Astrophys. J.* **654** (2007) 1006–1021
- [14] Scheck, L., Kifonidis, K., Janka, H.T., Müller, E.: Multidimensional supernova simulations with approximative neutrino transport. I. Neutron star kicks and the anisotropy of neutrino-driven explosions in two spatial dimensions. *Astron. Astrophys.* **457** (2006) 963–986
- [15] Scheck, L., Janka, H.T., Foglizzo, T., Kifonidis, K.: Multidimensional supernova simulations with approximative neutrino transport. II. Convection and the advective-acoustic cycle in the supernova core. *Astron. Astrophys.* **477** (2008) 931–952
- [16] Keil, W., Janka, H.T., Müller, E.: Ledoux Convection in Protoneutron Stars— A Clue to Supernova Nucleosynthesis? *Astrophys. J.* **473** (1996) L111
- [17] Burrows, A., Lattimer, J.M.: The birth of neutron stars. *Astrophys. J.* **307** (1986) 178–196
- [18] Pons, J.A., Reddy, S., Prakash, M., Lattimer, J.M., Miralles, J.A.: Evolution of Proto-Neutron Stars. *Astrophys. J.* **513** (1999) 780–804
- [19] Buras, R., Rampp, M., Janka, H.T., Kifonidis, K.: Two-dimensional hydrodynamic core-collapse supernova simulations with spectral neutrino transport. I. Numerical method and results for a $15M_{\odot}$ star. *Astron. Astrophys.* **447** (2006) 1049–1092
- [20] Müller, B., Janka, H., Dimmelmeier, H.: A New Multi-dimensional General Relativistic Neutrino Hydrodynamic Code for Core-collapse Supernovae. I. Method and Code Tests in Spherical Symmetry. *Astrophys. J. Suppl.* **189** (2010) 104–133
- [21] Rampp, M., Janka, H.T.: Radiation hydrodynamics with neutrinos. Variable Eddington factor method for core-collapse supernova simulations. *Astron. Astrophys.* **396** (2002) 361–392