

W. E. Nagel
M. Reschke

High Performance Computing in Science Engineering



High Performance Computing in Science and Engineering '11

Wolfgang E. Nagel • Dietmar B. Kröner •
Michael M. Resch
Editors

High Performance Computing in Science and Engineering '11

Transactions of the High Performance
Computing Center, Stuttgart (HLRS) 2011

 Springer

Editors

Wolfgang E. Nagel
Zentrum für Informationsdienste
und Hochleistungsrechnen (ZIH)
Technische Universität Dresden
Zellescher Weg 12-14
01069 Dresden
Germany
wolfgang.nagel@tu-dresden.de

Dietmar B. Kröner
Abteilung für Angewandte Mathematik
Universität Freiburg
Hermann-Herder-Str. 10
79104 Freiburg
Germany
dietmar@mathematik.uni-freiburg.de

Michael M. Resch
Höchstleistungsrechenzentrum
Stuttgart (HLRS)
Universität Stuttgart
Nobelstraße 19
70569 Stuttgart
Germany
resch@hlrs.de

Front cover figure: Laminar-turbulent transition on a common dolphin: Turbulent kinetic energy at 1 m/s and 1% turbulence intensity. (Simulation by D. Riedeberger and U. Rist, IAG, Stuttgart University, Germany. Geometry *DIXIE* kindly provided by V. Pavlov, FTZ, Kiel University, Germany)

ISBN 978-3-642-23868-0
DOI 10.1007/978-3-642-23869-7
Springer Heidelberg Dordrecht London New York

e-ISBN 978-3-642-23869-7

Library of Congress Control Number: 2011942929

Mathematics Subject Classification (2010): 65Cxx, 65C99, 68U20

© Springer-Verlag Berlin Heidelberg 2012

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Cover design: WMXDesign GmbH

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

It is a pleasure to announce that after the installation of the IBM Blue Gene/P at NIC/JSC in May 2009, as part of the Gauss Centre for Supercomputing (GCS) – linking together the three national supercomputing centres HLRS (Stuttgart), NIC/JSC (Jülich), and LRZ (Garching) – the next major step has been realized at HLRS with the installation of a new system, HERMIT, in October 2011. HERMIT is a large Cray XE-6 with a peak speed of more than 1 PFLOP/s. Based on the new AMD Interlagos chip, two sockets form a node, and HERMIT brings together 3,552 compute nodes with 113,664 cores, integrated into 38 water-cooled cabinets. Additionally, the system is tightly integrated with external servers for pre- and post-processing to support complex workflows. The new system entered the Top500 list in November 2011 and ranked number 12 on a worldwide level, achieving a Linpack value of 831,4 TFLOP/s. Having the new Petaflop system as an infrastructure, and together with the new research buildings for VISUS, SimTech, and HLRS, the Universität Stuttgart is well positioned to become one of the leading science nodes for simulation technology in Germany, as well as abroad.

In 2013, the second delivery phase will follow, and the final Cray system will then have a peak performance of roughly 5 PFLOP/s. Additionally, the LRZ will upgrade its own systems accordingly to a 3 PFLOP/s system in summer 2012. The plan is to have a Tier-0 HPC system within the GCS operating at any time within the five year period.

The HLRS also participates in the European project PRACE (Partnership for Advances Computing in Europe) as part of the GCS, extending its reach to all European member countries. Within the PRACE project, the GCS will provide access to high performance computing resources valued at 100 Million Euros. Moreover, the PRACE activities are well aligned with the HLRS activities in the European HPC support project, HPC-Europa2. Additionally, HLRS participates with partners in Germany in two Exascale Software Initiatives on European Level, namely TEXT and CRESTA where the challenges on the efficient use of current and future computing systems are investigated.

While the GCS has successfully addressed the high end computing needs, it was clear from the very beginning that an additional layer of support is required

to maintain the longevity of the Centre, via a network of competence centres across Germany. This gap is addressed by the Gauß–Allianz (GA), in which regional and local centres teamed up to create the necessary infrastructure, knowledge, and the required methods and tools. The mission of the Allianz is to coordinate the HPC-related activities of its members. By providing versatile computing architectures and by combining the expertise of the participating centres, the necessary ecosystem for computational science has been created. Strengthening the research and increasing the visibility to compete at the international level are further goals of the Gauß–Allianz. To disseminate information about its activities, the Gauß–Allianz has started to publish a flyer (GA-Infobrief, <http://www.gauss-allianz.de/infobrief>), issued several times a year.

A number of projects of the second BMBF HPC-call have started as early as April 2011. This call was directed towards proposals that enable and support petascale applications on more than 100,000 processors, as they are also currently available at HLRS. While the projects of the first funding round started in early 2009 and will complete within the next 6 months, the follow-up call had been delayed by more than 18 months. Nevertheless, all experts and administration authorities continue to acknowledge the strong need for such a funding program, given that the main issue identified in nearly all applications is that of *scalability*. The strategic funding plan involves another 20 Million Euros, with a yearly follow-up call over the next three years, for projects that develop scalable algorithms, methods, and tools to support massively parallel systems. This can be seen as a very large investment. Nevertheless, in relation to the investment in computing hardware within Germany over this five year period, the investment in software is still comparatively small, amounting to less than 20 per cent of the hardware investment. Furthermore, the investment in software will produce the ‘brains’ that will be needed to use the newly developed innovative methods and tools, to accomplish technological breakthroughs in scientific as well as industrial fields of applications.

It is widely known that the long term target is aimed not only at Petascale but at Exascale systems as well. We do not only need competitive hardware but also excellent software and methods to address – and solve – the most demanding problems in science and engineering. The success of this approach is of significant importance for our community, and will also greatly influence the development of new technologies and industrial products. Beyond being important, the success of this approach will finally determine whether Germany will be an accepted partner alongside the leading technology and research nations.

It is, therefore, a pleasure to announce that in October 2011, the German Research Foundation (DFG) has funded an additional Priority Program 1648 “Software for Exascale Computing (SPPEXA)” in the field of HPC. The funding is available for 6 years, starting January 2013 with 4 Million Euros per year, to support fundamental and basic research questions in several specific areas related to HPC.

Since 1996, the HLRS has supported the scientific community as part of its official mission. Just as in the past years, the major results of the last 12 months were presented at the 15th annual Results and Review Workshop on High Performance Computing in Science and Engineering, which was held on October 4–5, 2011 at

the Universität Stuttgart. The workshop proceedings contain the written versions of the research work presented. The papers were selected from all projects running at the HLRS and the SSC Karlsruhe during a one-year period between October 2010 and October 2011. Overall, a number of 47 papers were chosen from Physics, Solid State Physics, Chemistry, Reactive Flow, Computational Fluid Dynamics (CFD), Transport, Climate, and numerous other fields. The largest number of contributions originated from the CFD field, just as in many previous years, with 13 papers. Even though such a small collection cannot entirely represent an area this vast, the selected papers demonstrate the state-of-the-art in high performance computing in Germany. The authors were encouraged to emphasize the computational techniques used in solving the problems examined. This is an often forgotten aspect, and was the major focus of the workshop proceedings. Nevertheless, the importance of the newly computed scientific results for the specific disciplines is impressive.

We gratefully acknowledge the continuing support of the federal state of Baden-Württemberg in promoting and supporting high performance computing. Grateful acknowledgments are also due to German Research Foundation (Deutsche Forschungsgemeinschaft (DFG)) and the Germany Ministry for Research and Education (BMBF), as many projects pursued on the HLRS and SSC computing machines could not have been carried out without its support. Also, we thank Springer Verlag for publishing this volume and, thus, helping to position the national activities in an international framework. We hope that this series of publications contributes to the global promotion of high performance scientific computing.

Stuttgart, November 2011

Wolfgang E. Nagel
Dietmar Kröner
Michael Resch

Contents

Physics

P. Nielaba

The Influence of the Mass Ratio on Particle Acceleration by the Filamentation Instability	5
P. Kilian, T. Burkart, and F. Spanier	

The SuperN-Project: Neutrino Hydrodynamics Simulations of Core-Collapse Supernovae	15
B. Müller, L. Hudepohl, A. Marek, F. Hanke, and H.-Th. Janka	

Simulation of Pre-planetary Collisions with Smoothed Particle Hydrodynamics	29
R.J. Geretshauser, R. Speith, and W. Kley	

Copper Substrate Catalyzes Tetraazaperopyrene Polymerization	47
W.G. Schmidt, E. Rauls, U. Gerstmann, S. Sanna, M. Landmann, M. Rohrmüller, A. Riefer, S. Wippermann, and S. Blankenburg	

QCD Critical Surfaces at Real and Imaginary μ	57
O. Philipsen and Ph. de Forcrand	

Higgs Boson Mass Bounds from a Chirally Invariant Lattice Higgs-Yukawa Model	67
P. Gerhold, K. Jansen, and J. Kallarackal	

Massive and Massless Four-Loop Integrals	83
P. Baikov, K. Chetyrkin, J.H. Kühn, P. Marquard, and M. Steinhauser	

Solid State Physics

H. Fehske

Laser Ablation of Aluminium: Drops and Voids	93
J. Roth, J. Karlin, C. Ulrich, H.-R. Trebin, and S. Sonntag	

Cysteine on Gold: An <i>ab-initio</i> Investigation	105
B. Höffling, F. Ortmann, K. Hannewald, and F. Bechstedt	
Ab-initio Calculations of the Vibrational Properties of Nanostructures	119
G. Bester and P. Han	
Entropy and Metal-Insulator Transition in Atomic-Scale Wires: The Case of In-Si(111)(4×1)/(8×2)	131
W.G. Schmidt, E. Rauls, U. Gerstmann, S. Sanna, M. Landmann, M. Rohrmüller, A. Riefer, and S. Wippermann	
Obtaining the Full Counting Statistics of Correlated Nanostructures from Time Dependent Simulations	141
P. Schmitteckert	
Phase Diagram of the 1D t-J Model	153
A. Moreno, A. Muramatsu, and S. Manmana	
Chemistry	
C. van Wüllen	
Constrained Density Functional Theory of Molecular Dimers	169
J.-H. Franke, N.N. Nair, L. Chi, and H. Fuchs	
Atomistic Simulations of Electrolyte Solutions and Hydrogels with Explicit Solvent Models	185
J. Walter, S. Deublein, S. Reiser, M. Horsch, J. Vrabec, and H. Hasse	
cuVASP: A GPU-Accelerated Plane-Wave Electronic-Structure Code	201
S. Maintz, B. Eck, and R. Dronskowski	
Reacting Flows	
D. Kröner	
Assessment of Conventional Droplet Evaporation Models for Spray Flames	209
M.R.G. Zoby, A. Kronenburg, S. Navarro-Martinez, and A.J. Marquis	
Analysis of the Effects of Wall Boundary Conditions and Detailed Kinetics on the Simulation of a Gas Turbine Model Combustor Under Very Lean Conditions	229
F. Rebosio, A. Widenhorn, B. Noll, and M. Aigner	
Oxy-coal Combustion Modeling at Semi-industrial Scale	245
M. Müller, U. Schnell, and G. Scheffknecht	

Delayed Detached Eddy Simulations of Compressible Turbulent Mixing Layer and Detailed Performance Analysis of Scientific In-House Code TASCOM3D	259
M. Kindler, P. Gerlinger, and M. Aigner	
Computational Fluid Dynamics	
S. Wagner	
Discontinuous Galerkin for High Performance Computational Fluid Dynamics (hpcdg)	277
C. Altmann, A. Beck, A. Birkefeld, F. Hindenlang, M. Staudenmaier, G. Gassner, and C.-D. Munz	
Highly Efficient and Scalable Software for the Simulation of Turbulent Flows in Complex Geometries	289
D.F. Harlacher, S. Roller, F. Hindenlang, C.-D. Munz, T. Kraus, M. Fischer, K. Geurts, M. Meinke, T. Klühspies, V. Metsch, and K. Benkert	
A Computation Technique for Rigid Particle Flows in an Eulerian Framework Using the Multiphase DNS Code FS3D	309
P. Rauschenberger, J. Schlotke, and B. Weigand	
Optimization of Chaotic Micromixers Using Finite Time Lyapunov Exponents	325
A. Sarkar, A. Narváez, and J. Harting	
Numerical Simulation of Particle-Laden Turbulent Flows Using LES	337
M. Breuer and M. Alletto	
Large-Eddy Simulation of Supersonic Film Cooling at Finite Pressure Gradients	353
M. Konopka, M. Meinke, and W. Schröder	
Prediction of Stability Limits of Combustion Chambers with LES	371
B. Pritz, F. Magagnato, and M. Gabi	
Numerical Simulation of Laminar-Turbulent Transition on a Dolphin Using the γ-Re_θ Model	379
D. Riedeberger and U. Rist	
Wall Effects and Corner Separations for Subsonic and Transonic Flow Regimes	393
A. Klein, S. Illi, K. Nübler, T. Lutz, and E. Krämer	
Numerical Simulation of Helicopter Wake Evolution, Performance and Trim	409
F. Bensing, M. Embacher, M. Hollands, B. Kutz, M. Keßler, and E. Krämer	

Parameter Study for Scramjet Intake Concerning Wall Temperatures and Turbulence Modeling	425
B. Reinartz	
Unsteady Numerical Study of Wet Steam Flow in a Low Pressure Steam Turbine	437
J. Starzmann, M.V. Casey, and J.F. Mayer	
Turbulence Modelling for CFD-Methods for Containment Flows	451
A. Zirkel and E. Laurien	
Transport and Climate	
C. Kottmeier	
The Transport of Mineral Dust Towards Hurricane Helene (2006)	471
J. Schwendike, S. Jones, H. Vogel, and B. Vogel	
Numerical Modelling of Mediterranean Cyclones	489
C.-J. Lenz, U. Corsmeier, and C. Kottmeier	
Modelling Near Future Regional Climate Change for Germany and Africa	503
H.-J. Panitz, P. Berg, G. Schädler, and G. Fosser	
High-Resolution Climate Predictions and Short-Range Forecasts to Improve the Process Understanding and the Representation of Land-Surface Interactions in the WRF Model in Southwest Germany (WRFCLIM)	513
H.-S. Bauer, K. Warrach-Sagi, V. Wulfmeyer, T. Schwitalla, and M. Kim	
Direct Numerical Simulation and Implicit Large Eddy Simulation of Stratified Turbulence	523
S. Remmler and S. Hickel	
Miscellaneous Topics	
W. Schröder	
Allocation of Economic Capital in Banking: A Simulation Approach	541
H.-P. Burghof and J. Müller	
The Influence of Partial Melt on Mantle Convection	551
A.-C. Plesa and T. Spohn	
Molecular Modeling of Hydrogen Bonding Fluids: Phase Behavior of Industrial Fluids	567
S. Eckelsbach, M. Bernreuther, C. Engin, G. Guevara-Carrion, Y.-L. Huang, T. Merker, H. Hasse, and J. Vrabec	

“Brute-Force” Solution of Large-Scale Systems of Equations in a MPI-PBLAS-ScaLAPACK Environment	581
M. Roth, O. Baur, and W. Keller	
Metallic Foam Structures, Dendrites and Implementation Optimizations for Phase-Field Modeling	595
A. Vondrous, B. Nestler, A. August, E. Wesner, A. Choudhury, and J. Hötzer	
Quaero 2010 Speech-to-Text Evaluation Systems	607
S. Stüker, K. Kilgour, and F. Kraft	
Accurate Simulation of Wireless Vehicular Networks Based on Ray Tracing and Physical Layer Simulation	619
T. Gaugel, L. Reichardt, J. Mittag, T. Zwick, and H. Hartenstein	
Reduction of Numerical Sensitivities in Crash Simulations on HPC-Computers (HPC-10)	631
O. Mangold, R. Prohl, A. Tkachuk, and V. Trickov	
Three-Dimensional Gyrotron Simulation Using a High-Order Particle-in-Cell Method	637
A. Stock, J. Neudorfer, B. Steinbusch, T. Stindl, R. Schneider, S. Roller, C.-D. Munz, and M. Auweter-Kurtz	

Physics

Prof. Dr. Peter Nielaba

Many important results have been achieved by the computer time granted at the HLRS. The contributions in the proceedings present the results of large scale simulations for astrophysics phenomena, nano-systems, and elementary particle models and are summarized and commented below.

Patrick Kilian, Thomas Burkart, and Felix Spanier (University of Würzburg, “*iotmrofi*”) have studied the influence of the mass ratio on particle acceleration by the filamentation instability. Observations indicate that several types of astrophysical sources produce relativistic jets that interact with the intergalactic medium, creating regions of counter-streaming plasma. Under these conditions the plasma is susceptible to filamentation instabilities. The analytical analysis of this environment is highly non-trivial, which leads to the extensive use of computer simulations to study these conditions and the connection to the energetic photons and particles emanating from these sources. To make simulations feasible one has to make a couple of simplifications to reduce the computational complexity to a level that is reachable with today's computers. One such simplification is the reduction of the proton mass compared to the electron mass. The authors try to assess what the lower limit of this quantity is that still allows a realistic representation of the situation in nature. The authors have used the particle-in-cell code ACRONYM, in parts on the NEC Nehalem Cluster at the High Performance Computing Center Stuttgart (HLRS). The largest jobs there used 64 nodes with 8 processors per core, the largest simulation used 37709 CPU hours and was split into 17 consecutive jobs using 128 CPUs each. In that time $3.4 \cdot 10^9$ particles were moved through 3600 time steps.

B. Müller, L. Hudepohl, A. Marek, F. Hanke, and H.-Th. Janka from the MPI for Astrophysics in Garching (“*SuperN*”) have investigated two-dimensional (core collapse) supernova models by simulations, give an overview on the relevant equations and the algorithm for its solution that are employed in their code, and report on their efforts to improve the physics in their supernova code VERTEX as well as its computational efficiency. Recent results of simulations performed on the NEC SX-8

at the HLRS include findings about the role of general relativity in core-collapse supernovae, about nucleosynthesis conditions in O-Ne-Mg core supernovae, and about the proto-neutron star cooling phase.

R.J. Geretshauser, R. Speith, F. Meru, and W. Kley (University of Tübingen, “SPH-PPC”) have analyzed pre-planetesimal collisions with their solid body—smoothed particle hydrodynamics (SPH) code *parasph*. The formation of planetesimals requires the right amount of sticking, bouncing, and fragmentation to be consistent with observations. Therefore, collisions of pre-planetesimals have to be investigated as thoroughly as possible and their outcome has to be mapped as precisely as possible depending on all relevant parameters such as initial porosity, collision partner size, and impact velocity. The code *parasph* is based on the ParaSPH library, featuring domain decomposition, load balancing, nearest neighbor search, and inter-node communication, extended for the simulation of elasticity and plasticity including the time evolution of the deviatoric stress tensor and by an implementation of a porosity model. The parallel implementation utilizes the Message Passing Interface (MPI) library, and HDF5 was included as a compressed input and output file format with increased accuracy, decreasing the amount of required storage space. The authors have developed the first damage model which is based on the inhomogeneity of SiO₂ dust aggregates, and they have investigated the occurrence of sticking and bouncing for macroscopic and microscopic aggregates and head-on collisions of pre-planetesimals. The simulations were carried out on the NEC Nehalem cluster of the HLRS with 240,143 to 476,476 SPH particles depending on the size of the projectile. 32 to 80 cores were used, and simulations roughly took 72 to 240 h for 1 s of simulated time, depending on size of the problem and involved physical process.

W.G. Schmidt, E. Rauls, U. Gerstmann, S. Sanna, M. Landmann, M. Rohrmüller, A. Riefer, S. Wippermann, S. Blankenburg (University of Paderborn, “MolArchI”) have investigated by density functional theory (DFT) calculations the polymerization of 1,3,8,10-tetraazaperopyrene (TAPP) molecules on a Cu(111) substrate, as observed in recent STM experiments. According to the authors computations, the substrate catalyzes this tautomerization, and metal coordinated (provided the process is accompanied by a dehydrogenation) as well as uncoordinated polymerization processes are possible. The authors conclude, that the catalytic effect of metallic substrates may thus assist in the formation of covalently bonded molecular networks which are not formed in the gas phase or in solution. In the calculations, the electron-ion interaction was described by the projector-augmented wave (PAW) method with a relatively moderate energy cutoff of 340 eV, and the adstructures were modeled in periodically repeated supercells, containing two atomic Cu layers. The calculations within this project were performed on the NEC SX-8 and SX-9 of the Höchstleistungsrechenzentrum Stuttgart. The DFT calculations were performed within the local density approximation (LDA) for exchange and correlation as implemented in VASP. The ground-state DFT calculations have been parallelized for different bands and sampling points in the Brillouin zone using the message passing interface (MPI), and parallelization over bands and plane wave coefficients at the same time reduced the communication overhead.

Ph. de Forcrand from the ETH Zürich and O. Philipsen from the University of Münster (“*muQCD*”) have calculated the critical surface bounding the region featuring chiral phase transitions in the quark mass and chemical potential parameter space of quantum chromo dynamics (QCD) with three flavors of quarks. Their calculations are valid for small to moderate quark chemical potentials, $\mu \leq T$. Previous calculations were done on coarse $N_t = 4$ lattices, corresponding to a lattice spacing of a ~ 0.3 fm. Now, the authors present updated results for three degenerate flavors at zero and finite density on $N_t = 6$ lattices, corresponding to a lattice spacing of a ~ 0.2 fm, and for the phase structure at imaginary chemical potential $\mu/T = i\pi/3$, finding tricritical lines which bound the continuation of the chiral as well as the deconfinement transition surfaces to imaginary chemical potentials, and explain their curvature. For their Monte Carlo simulations the authors used the standard Wilson gauge and Kogut-Susskind fermion actions. Configurations are generated using the Rational Hybrid Monte Carlo (RHMC) algorithm. In order to investigate the critical behavior of the theory, the authors use the Binder cumulant as an observable. For each set of fixed quark mass and chemical potential, the critical coupling β_c has been interpolated from a range of typically 3–5 simulated β -values by Ferrenberg-Swendsen reweighting. The simulations have been performed on the NEC SX-8 at the HLRS in Stuttgart and the EEE Grid at CERN. An estimate of the Binder cumulant for one set of mass values consisted of at least 200k trajectories, and the estimate of a critical point required at least 500k trajectories. The derivatives on the $N_t = 4$ lattices are based on 500k trajectories, on $N_t = 6$ the authors have so far also collected 1M trajectories.

Philipp Gerhold, Karl Jansen, and Jim Kallarackal from the Humboldt University and DESY Zeuthen (“*Xyukawa*”) considered a chirally invariant lattice Higgs-Yukawa model based on the Neuberger overlap operator. The model has been evaluated using PHMC-simulations, and the authors present final results on the upper and lower Higgs boson mass bound. The question of a fourth generation of heavy quarks has recently gained attention, and the authors illustrate preliminary results of the Higgs boson mass bounds within this framework. The authors as well discuss their progress on properties of the Higgs boson with respect to its unstable nature, such as the decay width and the resonance mass of the Higgs boson. The current program, which is running on the XC2 (Karlsruhe), has been improved in many ways, including algorithmic as well as technical issues. The authors detected, that most of the used computer time is spent on performing the Fast Fourier Transforms. The implemented parallelization techniques therefore mainly focused on exploiting the available multi processor system for calculating the FFT, mostly limited by memory access speeds. The authors idea was to align processes to memory segments where the memory had been allocated, leading to a good scaling behavior. At present, the authors report on runs on 32^4 and 40^4 lattices on the *fat* nodes having 8 cores each. The simulation program will need approximately 20 GB of main memory on each node for the 40^4 -lattice, and the fat nodes with their 128 GB of memory are therefore well suited for the computations. About 200 simulations are stored in archive space and occupy roughly 8.4 TB.

P. Baikov, K. Chetyrkin, J.H. Kühn, P. Marquard and M. Steinhauser from the KIT Karlsruhe (“*ParFORM*”) have investigated massive and massless four-loop integrals, the computations were mainly performed on the Landeshöchstleistungsrechner XC4000. The problems treated within their project aim for the evaluation of so-called Feynman diagrams which in turn lead to quantum corrections within a given quantum field theory like Quantum Electrodynamics or Quantum Chromodynamics but also supersymmetric theories. Each Feynman diagram has a one-to-one translation to a high-dimensional momentum integral to be computed, if possible analytically. There are several algorithms which have been suggested for the computation. All of them require huge resources of computer power which can only be handled in combination with effective programs. The workhorse of the authors for such calculations is the computer algebra program FORM and its parallel versions ParFORM and TFORM, developed at the authors institute. Since August 2010 all versions of FORM are open source. The parallelization concept for FORM is roughly described as follows: The original expression is divided into several pieces which are then distributed to the individual processors or cores (workers). Once the workers have finished their job the resulting expressions have to be collected by one processor which combines the results. A computer architecture running ParFORM or TFORM requires a fast connection to the (in general) local hard disks of the order of one tera byte per core. The typical CPU time reaches from several hours to several months depending on the concrete problem under consideration. One part of the authors project deals with two different kind of integrals. The authors investigated massive vacuum diagrams, also called tadpoles. These diagrams play an important role in the determination of the low-energy expansion of the vacuum polarization functions which can be used for the extraction of the masses of the charm and bottom quark from experimental data. The calculation of the low-energy expansion and the analysis of the available data has led to the most precise determination of the masses of the heavy quarks. A new class of integrals, so-called four-loop onshell integrals, appear as building blocks in important physics applications like the $\overline{\text{MS}}$ -onshell relation or the anomalous magnetic moment of the muon. The $\overline{\text{MS}}$ -onshell relation allows to relate the values of the quark masses in different renormalization schemes and is up to now only known at three-loop order. First results are already available: the calculation of the contributions from diagrams with at least two closed massless quark loops. The calculation of the anomalous magnetic moment of the muon, which, together with the one of the electron, is one of the most precisely measured quantities, is scheduled by the authors, once the calculation of the $\overline{\text{MS}}$ -onshell relations is complete. Since it is based on the same families of integrals as the $\overline{\text{MS}}$ -onshell relation many results can be reused.

The Influence of the Mass Ratio on Particle Acceleration by the Filamentation Instability

Patrick Kilian, Thomas Burkart, and Felix Spanier

Abstract Observations indicate that several types of astrophysical sources produce relativistic jets that interact with the intergalactic medium, creating regions of counterstreaming plasma. Under these conditions the plasma is susceptible to filamentation instabilities. Analytical analysis of this environment is highly non-trivial, which leads to the extensive use of computer simulations to study these conditions and the connection to the energetic photons and particles emanating from these sources. To make simulations feasible one has to make a couple of simplifications to reduce the computational complexity to a level that is reachable with todays computers. One such simplification is the reduction of the proton mass compared to the electron mass. This project tries to assess what the lower limit of this quantity is that still allows a realistic representation of the situation in nature.

1 Introduction

In this project we wanted to study the influence of the mass ratio protons vs electrons on the particle acceleration due to the filamentation instability in relativistic plasma. In nature we find relativistic counterstreaming plasmas in interaction regions between the intergalactic medium and jets from sources like active galactic nuclei (AGNs) and gamma-ray bursts (GRBs). These sources are—just like supernova remnants (SNRs)—known to produce more energetic photons than expected from a purely thermal source [1]. This non-thermal tail extends to extremely large energies and often shows a connection between energy and flux following a power-law. Furthermore we can infer from observations of strong synchrotron radiation that these sources also generate accelerate particles up to very high energies. Unlike the photon spectrum the particle spectrum is not directly observable here on earth, but most likely it follows a power-law too.

Patrick Kilian · Thomas Burkart · Felix Spanier
Institut für Theoretische Physik und Astrophysik, Universität Würzburg, Emil-Fischer-Str. 31,
97074 Würzburg, Germany, e-mail: fspanier@physik.uni-wuerzburg.de

The precise mechanism that accelerates particles up to the Petaelectronvolt scale which we observe in particles reaching earth as part of the cosmic rays is not entirely clear yet. Several plausible mechanisms including Fermi acceleration of type I or II are known from theoretical considerations and semi-analytical calculations. AGNs and GRBs with their relativistic outflows allow also for a different mechanism, the filamentation instability. This plasma instability may occur whenever collisionless plasmas are counterstreaming in an unmagnetized medium. Here not only magnetic fields are created, but also electrical fields may accelerate particles. The filamentation instability is especially interesting as it may provide the necessary pre-acceleration required for the Fermi-I mechanism.

The non-isotropic and non-homogeneous situation in the jet, the non-linear coupling between particles and electromagnetic fields and the collective and correlated behavior of the plasma itself require either quite extensive simplifying assumptions to make analytical calculations feasible or self-consistent three-dimensional simulations. The tool of choice for self-consistent simulations of collisionless plasmas are Particle-in-cell codes. For the simulations our in-house code ACRONYM was used that is described in Sect. 3.

Even simulations on the largest available computers can't track position and momentum of each and every single particle in a relativistic jet. So even large-scale simulations need some albeit less drastic simplifications. The first simplification one usually makes is lumping a number of particles of one species, say one billion electrons, which are close together in phase space to create one metaparticle with accordingly larger charge and mass. This leaves the q/m ratio constant and thus doesn't alter the acceleration caused by the Lorentz force. Obviously there is a limit how many particles may be lumped together and how many metaparticles need to remain which will be discussed later on.

The next step is to replace the direct interaction between particles due to the coulomb force by an indirect interaction via the strength of the electromagnetic field, stored on a spatially discrete grid. Using this particle-mesh method improves the computational complexity from $\mathcal{O}(n^2)$ to $\mathcal{O}(n)$ in the particle number. This method underrepresents short range forces (between particles that are much closer to each other than the size of a grid cell) but reflects long range forces quite accurately [2]. In the environments mentioned above long range interactions dominate rendering short range interactions comparatively unimportant, hence the plasma is often referred to as being "collisionless".

The third big simplification concerns the mass ratio between protons and electrons and is the topic of this project. The rate at which protons evolve in time and form filaments is closely connected to the mass of the proton m_p . However to faithfully represent the plasma the simulation code has to resolve the short-range and transient fluctuations which are mainly carried by electrons and are therefore connected to the electron mass m_e . The typical length scale of these fluctuations is the Debye length λ_D [3] which restricts the edge length of the grid cells to $dx < \lambda_D$ [4]. In turn the Courant-Friedrichs-Lewy condition [5] limits the length of each explicit time step to $dt < 3^{-1/2}c/\Delta$. If one wants to look at the behavior at late times, say after a couple of proton gyro periods, the simulation has to run for ten thousands of

time steps. To alleviate this problem one can reduce the mass of the protons in the simulation, creating mass ratios smaller than the rather large value of 1836 that we find in nature. So in other words this project attempts to shed light on the question “How small can we make the mass ratio without misrepresenting nature”.

2 Scientific Results

The description of the project status closely follows the PhD thesis of Thomas Burkart [6] who ran and analyzed most of the simulations presented here. A more detailed scientific treatment of the results can be found in [7].

For numerical reasons the simulations were conducted neither in the rest frame of the background medium nor in the rest frame of the jet plasma but rather in the frame where the common center of mass is at rest. In this frame both populations are streaming with a velocity close to the speed of light.

Initially the density of particles is uniform in the plane normal to the streaming direction for both populations. As the filamentation instability grows it creates strong flux tubes in which particles from one initial population (either jet or background) propagate preferentially creating a net current in the flux tube. This current generates a strong magnetic field around the flux tube, further compressing it. The magnetic field around the flux tubes has large components in the plane normal to the streaming direction. Thus the energy stored in the two perpendicular magnetic field components are a good proxy for the growth of the flux tubes.

Figure 1 shows five reconstructed field lines of the magnetic field in the neighborhood of a flux tube. The flux tube itself is represented by the particle density shown semi-transparently in gray-scale. The component of the magnetic field in the direction of the flux tube is rather small and the field lines form nearly closed loops around the flux tube, compressing it.

2.1 *Small and Intermediate Mass Ratios*

Coming back to the topic of the mass ratio it is obvious that the light electrons react much stronger to a given magnetic field than the heavy protons. Therefore small fluctuations in the electron density suffice to generate enough magnetic field to bunch the electrons a little closer, amplifying the initial fluctuation.

If the mass of a proton is not much larger than the electron mass, the protons will be strongly affected by the small magnetic fields too. This has two effects: The first is that the protons are—due to their opposite charge—repulsed from the magnetic field that bundles the electrons. As the protons are ejected from the forming electron flux tubes, charge separation increases which amplifies the magnetic field. The expelled protons form flux tubes too, which tighten due to the same filamentation instability, creating proton flux tubes on the same timescale as the electron flux




Fig. 1 Exemplary reconstructed field lines (shown in green) around a single flux tube stay almost completely in the perpendicular plane. The background image in gray-scale shows the mass density which peaks in the flux tubes

tubes, much faster than one would expect based on the mass ratio. The second effect is that the instability of the electrons happens slower than in the case of the physical mass ratio. This is due to the fact that moving charged particles within a magnetic field create an opposite magnetic field following Lenz' rule. As the magnetic field that was originally created by the electrons pushes the protons away it is weakened and the bunching of the electrons is slowed down.

Both effects can be seen in Fig. 2. For all mass ratios shown in this plot electron and proton instability generate one shared peak of the perpendicular field and larger mass ratios shift that peak to larger times. Mass ratios 1 and 5 show similar results, the differences are no larger than the statistical fluctuations one would find for two runs of equal mass ratio. A mass ratio of 20 that is often used in kinetic plasma simulations results in a slightly later peak after $60 \omega_{pe}^{-1}$ instead of $44 \omega_{pe}^{-1}$. A mass ratio of 42.8 (the square root of the mass ratio found in nature) results in an even later




Fig. 2 Temporal development of the perpendicular magnetic field for small and intermediate mass ratios

($68.3 \omega_{pe}^{-1}$) peak. Furthermore two distinct growth periods from $20 \omega_{pe}^{-1}$ to $35 \omega_{pe}^{-1}$ and from $55 \omega_{pe}^{-1}$ to $68 \omega_{pe}^{-1}$ can be seen but the separation between the two is not large enough to allow the electron flux tubes to decay before the proton flux tubes develop.

2.2 High Mass Ratios

For mass ratios larger than the ones discussed in the subsection above the electron and proton instability decouple and happen on their own undisturbed timescales. This can be seen quite well in Fig. 3 where the electron instability happens in the range from $30 \omega_{pe}^{-1}$ to $50 \omega_{pe}^{-1}$ independently from the mass of the protons.

The time of the peak magnetic field of the proton instability follows the linear relation $t_{peak} = 33.4 \omega_{pe}^{-1} + 0.8 \omega_{pe}^{-1} m_p/m_e$ very closely. Without wanting to over interpret this linear fit through just three data point we claim that the time scale of the proton instability grows linearly with the proton mass was expected from theory.

Given that the electron instabilities decay to half of their peak value by the time $57 \omega_{pe}^{-1}$ have passed and that the proton instability needs—just like the electron instability—at least $20 \omega_{pe}^{-1}$ to grow one can conclude that a mass ratio of 42.8 is indeed too small but 100 suffice to decouple the two fairly well.




Fig. 3 Temporal development of the perpendicular magnetic field for large mass ratios

2.3 Other Influences

To check if a mass ratio of 100 is indeed large enough as claimed in the preceding subsection we ran two more simulations. Both use the same mass ratio but are changed in either composition or size of the simulation box.

Looking at Fig. 4 the dashed line belongs to the simulation where both plasma populations are hydrogen plasma. Contrast this with the other simulations where one direction contains fewer protons and some positron (so effectively a mixture of hydrogen and pair plasma). The smaller number of light constituents results in a smaller than usual electron peak and the larger number of protons enhances the second peak in the magnetic field. The timing however remains unaffected.

The other investigated variant is shown as a dotted line in Fig. 4. In this case the simulations used twice the number of grid cells in both perpendicular directions, increasing the size of the simulation box in physical units correspondingly. The time evolution up to the point of $60 \omega_{pe}^{-1}$ which is dominated by the electron instability remains nearly unchanged. However the following proton instability grows slower and longer, reaching a later and stronger peak in the perpendicular magnetic field. The most likely reason is that a larger simulation domain reduces interaction with other filaments and more importantly the interaction of filaments with their own mirror image via the periodic boundaries. This point is the topic of further studies but is outside the focus of this project and doesn't affect the key conclusion that a




Fig. 4 Influences on the temporal development of the perpendicular magnetic field other than the mass ratio

mass ratio significantly smaller than 100 results in an unphysical coupling between electrons and protons.

3 Numerical Performance

The particle-in-cell code ACRONYM that has been developed at the chair of astronomy at the University of Würzburg has been described in detail in the application of this project. The following paragraph gives a short overview of the code and the main features for reader who are not familiar with the code. A particle-in-cell code belongs to the class of particle mesh methods commonly used in numerical studies of many-body problems. For a detailed introduction see [2] or [8]. The field quantities like electric and magnetic fields as well as current densities are stored in a Yee lattice [9]. The particles deposit currents on the grid in each time step using Esirkepov's method [10] which influences the electric and magnetic fields in the next update step. The electromagnetic fields act on the particles through the Lorentz force which is implemented using the Boris push [11]. A more detailed albeit slightly dated description of our code can be found in [12].

Some of the simulations of this project were run in-house, the other were done using the NEC Nehalem Cluster at the High Performance Computing Center Stuttgart (HLRS). The largest jobs there used 64 nodes with 8 processors per node. Our code

would easily scale to the 512 cores reserved this way but we had some problems with MPI threads running out of memory. Consequently most runs were conducted using two or four MPI threads per node to increase the amount of RAM available to each thread. The source of the problem is that a huge amount of particles is concentrated inside the flux tube, resulting in quite unbalanced memory requirements between the MPI threads depending on the presence or absence of a large flux tube in the part of the simulation domain covered by each MPI thread. As this was the first time that the available memory was the limiting factor our code has no support for OpenMP yet. Currently we are working on a dynamic rebalancing of the computational domains between different MPI threads which should remove the problem of the imbalanced memory usage. This will remove the need to run with a reduced number of threads per node and consequently the incentive to add OpenMP support will likely go away again.

The largest simulation in this project used 37709 CPU hours and was split into 17 consecutive jobs using 128 CPUs each. In that time $3.4 \cdot 10^9$ particles were moved through 3600 time steps. This is just short of 700 particle updates per CPU and second. This falls well short of the peak performance of 200000 particles per second of our code. The reason is that each job had to read 250 gigabytes of particles as well as several gigabytes of electromagnetic fields from disk to reconstruct the state that the preceding job reached as well as write back all that information at the end of the run. The majority of the 15 to 20 hours runtime was spend waiting for the IO to happen. Other jobs which used one fifth the number of particles per cell where not impeded by this problem and performed as expected. Recent improvements in the IO part of the code have shown a large speed up on other systems and we plan to port the improvement to the code version running at HLRS prior to any new simulations.

References

1. F.A. Aharonian, A.G. Akhperjanian, K.M. Aye, A.R. Bazer-Bachi, M. Beilicke, W. Benbow, D. Berge, P. Berghaus, K. Bernlöhr, O. Bolz, C. Boisson, C. Borgmeier, F. Breitling, A.M. Brown, J. Bussons Gordo, P.M. Chadwick, V.R. Chitnis, L.M. Chounet, R. Cornils, L. Costamante, B. Degrange, A. Djannati-Ataï, L.O. Drury, T. Ergin, P. Espigat, F. Feinstein, P. Fleury, G. Fontaine, S. Funk, Y.A. Gallant, B. Giebels, S. Gillesen, P. Goret, J. Guy, C. Hadjichristidis, M. Hauser, G. Heinzelmann, G. Henri, G. Hermann, J.A. Hinton, W. Hofmann, M. Holleran, D. Horns, O.C. de Jager, I. Jung, B. Khélifi, N. Komin, A. Konopelko, I.J. Latham, R. Le Gallou, M. Lemoine, A. Lemièrre, N. Leroy, T. Lohse, A. Marcowith, C. Masterson, T.J.L. McComb, M. de Naurois, S.J. Nolan, A. Noutsos, K.J. Orford, J.L. Osborne, M. Ouchrif, M. Panter, G. Pelletier, S. Pita, M. Pohl, G. Pühlhofer, M. Punch, B.C. Raubenheimer, M. Raue, J. Raux, S.M. Rayner, I. Redondo, A. Reimer, O. Reimer, J. Ripken, M. Rivoal, L. Rob, L. Rolland, G. Rowell, V. Sahakian, L. Saugé, S. Schlenker, R. Schlickeiser, C. Schuster, U. Schwanke, M. Siewert, H. Sol, R. Steenkamp, C. Stegmann, J.P. Tavernet, C.G. Théoret, M. Tluczykont, D.J. van der Walt, G. Vasileiadis, P. Vincent, B. Visser, H.J. Völk, S.J. Wagner, *Nature* **432**, 75 (2004). DOI [10.1038/nature02960](https://doi.org/10.1038/nature02960)
2. R.W. Hockney, J.W. Eastwood, *Computer simulation using particles* (Bristol: Hilger, 1988)
3. Mitteilungen der Astronomischen Gesellschaft Hamburg **13**, 29 (1959)

4. B.I. Cohen, A.B. Langdon, D.W. Hewett, R.J. Procassini, *Journal of Computational Physics* **81**(1), 151 (1989). DOI [10.1016/0021-9991\(89\)90068-5](https://doi.org/10.1016/0021-9991(89)90068-5)
5. R. Courant, K. Friedrichs, H. Lewy, *Mathematische Annalen* **100**(1), 32 (1928). DOI [10.1007/BF01448839](https://doi.org/10.1007/BF01448839)
6. T. Burkart, *Der Einfluss des fundamentalen Massenverhältnisses auf die Teilchenbeschleunigung durch Plasmainstabilitäten*. Ph.D. thesis, Julius-Maximilians-Universität Würzburg (2010)
7. T. Burkart, O. Elbracht, U. Ganse, F. Spanier, *ApJ* **720**, 1318 (2010). DOI [10.1088/0004-637X/720/2/1318](https://doi.org/10.1088/0004-637X/720/2/1318)
8. C.K. Birdsall, A.B. Langdon, *Plasma physics via computer simulation*, 1st edn. (New York: Taylor and Francis, 2005)
9. K. Yee, *Antennas and Propagation*, *IEEE Transactions on* **14**(3), 302 (1966)
10. T.Z. Esirkepov, *Computer Physics Communications* **135**(2), 144 (2001). DOI [10.1016/S0010-4655\(00\)00228-9](https://doi.org/10.1016/S0010-4655(00)00228-9)
11. G. Penn, P.H. Stoltz, J.R. Cary, J. Wurtele, *Journal of Physics G: Nuclear and Particle Physics* **29**(8), 1719 (2003). URL <http://stacks.iop.org/0954-3899/29/1719>
12. T. Burkart, O. Elbracht, F. Spanier, *Astronomische Nachrichten* **328**, 662 (2007)

The SuperN-Project: Neutrino Hydrodynamics Simulations of Core-Collapse Supernovae

B. Müller, L. Hudepohl, A. Marek, F. Hanke, and H.-Th. Janka

Abstract We give an overview of the challenges and the current status of our two-dimensional (core collapse) supernova modelling, and present the system of equations and the algorithm for its solution that are employed in our code VERTEX. We also discuss the parallelisation of VERTEX, give scaling results on different architectures, and report on our ongoing efforts to increase the computational efficiency of the code. Furthermore, we outline some of the recent results obtained from simulations performed on the NEC SX-8 at the HLRS. Specifically, we report our findings about the role of general relativity in core-collapse supernovae, about nucleosynthesis conditions in O-Ne-Mg core supernovae, and about the proto-neutron star cooling phase.

1 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

B. Müller · L. Hudepohl · A. Marek · F. Hanke · H.-Th. Janka
Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Strasse 1, Postfach 1317, D-85741
Garching bei München, Germany, e-mail: bjmuellr@mpa-garching.mpg.de

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?

2 Numerical Models

2.1 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 reenergise 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., [3–5]), 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 [7–9].
- (2) Recent work [10–13] 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 [16–18] 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-magnetised 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.

2.2 *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 parallelisation 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.




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

2.3 “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, $f(r, \vartheta, \phi, \varepsilon, \Theta, \Phi, t)$. Equivalently, the neutrino intensity $I = c/(2\pi\hbar c)^3 \cdot \varepsilon^3 f$ 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 discretisation (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

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

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.

This leaves us with a four-dimensional problem. So far no approximations have been made. In order to reduce the size of the problem even further, one needs to