W. E. Nagel · D. Kröner
M. Resch (Eds.)

High Performance Computing in Science and Engineering '07

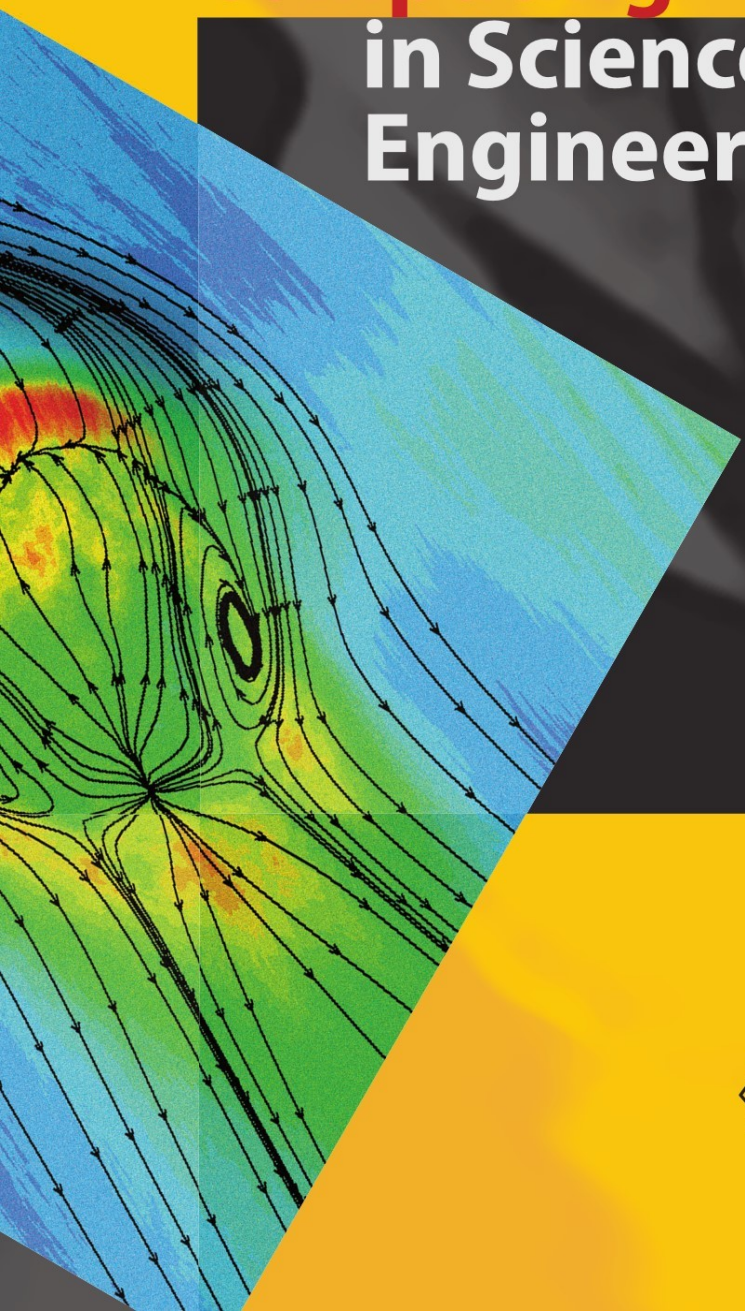

HLRS

Springer

High Performance Computing in Science and Engineering ’07

Wolfgang E. Nagel · Dietmar Kröner · Michael Resch
Editors

High Performance Computing in Science and Engineering ’07

Transactions of the High Performance Computing Center Stuttgart (HLRS) 2007

Springer

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

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

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

Front cover figure: Streamlines of the mean flow and turbulent kinetic energy contours obtained from a simulation of separated flow over an axisymmetric hill, Institute for Hydromechanics, University of Karlsruhe

ISBN 978-3-540-74738-3 e-ISBN 978-3-540-74739-0

DOI 10.1007/978-3-540-74739-0

Library of Congress Control Number: 2007939457

Mathematics Subject Classification (2000): 65Cxx, 65C99, 68U20

Typesetting: by the editors using a Springer TEX macro package
Production and data conversion: LE-TEX Jelonek, Schmidt & Vöckler GbR, Leipzig, Germany
Cover design: WMXDesign, Heidelberg

Printed on acid-free paper

9 8 7 6 5 4 3 2 1

springer.com

Preface

Again, the last year has been very successful for high performance computing in Baden-Württemberg and beyond. The NEC SX-8 vector supercomputer is highly utilized and has been very successfully used by many scientific projects. In February 2007, the new "Landes-Höchstleistungsrechner HP XC4000" has been inaugurated at the SSC in Karlsruhe. This massively parallel high performance system complex, built from hundreds of Intel Itanium processors and more than three thousand AMD Opteron cores, has helped in doing additional science on a very high level. Nevertheless, the time has already come to start thinking about the next follow-up system in future.

The discussions and plans on all scientific, advisory, and political levels to realize an even larger "European Supercomputer" in Germany, where the hardware costs alone will be around 200 Million Euro, are getting closer to realization. There are many good reasons to invest in such a program because – beyond the infrastructure – such a scientific research tool will attract the best brains to tackle the problems related to the software and methodology challenges. As part of the strategy, the three national supercomputing centres HLRS (Stuttgart), NIC/JSC (Jülich) and LRZ (Munich) have formed the Gauss-Centre for Supercomputing (GCS) as a new virtual organization enabled by an agreement between the Federal Ministry of Education and Research (BMBF) and the state ministries for research of Baden-Württemberg, Bayern, and Nordrhein-Westfalen from July 2006. Already today, the GCS provides the most powerful high-performance computing infrastructure in Europe.

Moreover, it is expected that in the next few months – following the proposal of the German HPC community, guided by Professor Andreas Reuter (EML) – the reshape of the High Performance Computing in Germany will proceed to form the German HPC "Gauss Alliance", with the goal to improve and establish competitiveness for the coming years. Beyond the stabilization and strengthening of the existing German infrastructures – including the necessary hardware at a worldwide competitive level – a major software research and support program to enable Computational Science and Engineering on the

required level of expertise and performance – which means: running Petascale applications on more than 100,000 processors – is promised by the BMBF. It is expected that for the next years 20 Million Euro will be spend – on a yearly basis – for projects to develop scalable algorithms, methods and tools in many areas to support massively parallel systems. As we all know, we do not only need competitive hardware but also excellent software and methods to approach – and solve – the most demanding problems in science and engineering. To achieve this challenging goal every three-year project supported by that program will need to integrate excellent research groups at the universities with colleagues from the competence network of HPC centers in Germany. The success of this approach is of utmost importance for our community and also will strongly influence the development of new technologies and industrial products; beyond that, this will finally determine if Germany will be an accepted partner among the leading technology and research nations.

The role of national supercomputing centers like HLRS can never be limited to the pure operation and services of hardware systems. Research in methods and tools is strongly necessary to be able to support users in exploiting the full potential of supercomputing systems. HLRS has emphasized its leading role in this field over the last years. Most recently, as part of the German national elite research program HLRS and its research partners at the University of Stuttgart were awarded a funding of 30 Million Euro for the next five years for a cluster of excellence in Simulation Technology. With this success, the University of Stuttgart and HLRS will further strengthen their national and international position as centers for high performance computing and simulation technology.

The goal of the cluster of excellence in simulation technology is to go from isolated approaches (numerical kernels, methods, and tools) to an integrated system science. Research areas will include molecular and particle simulation, advanced mechanics of multi-scale and multi-field problems, systems analysis and inverse problems, numerical and computational mathematics, integrated data management and interactive visualization, hybrid high-performance computing systems and simulation software engineering and an integrative platform of reflection and evaluation.

Since 1996, HLRS is supporting the scientific community as part of its official mission. Like in the years before, the major results of the last 12 months were reported at the Tenth Results and Review Workshop on High Performance Computing in Science and Engineering, which was held October 4–5, 2007 at Stuttgart University. This volume contains the written versions of the research work presented. The papers have been selected from all projects running at HLRS and at SSC Karlsruhe during the one year period beginning October 2006. Overall, 43 papers have been chosen from Physics, Solid State Physics, Computational Fluid Dynamics, Chemistry, and other topics. The largest number of contributions, as in many other years, came from CFD with 17 papers. Although such a small collection cannot represent a large area in total, the selected papers demonstrate the state of the art in high perfor-

mance computing in Germany. The authors were encouraged to emphasize computational techniques used in solving the problems examined. This often forgotten aspect was the major focus of this volume, nevertheless this should not disregard the importance of the newly computed scientific results for the specific disciplines.

We gratefully acknowledge the continued support of the Land Baden-Württemberg in promoting and supporting high performance computing. Grateful acknowledgement is also due to the Deutsche Forschungsgemeinschaft (DFG): many projects processed on the machines of HLRS and SSC could not have been carried out without the support of the DFG. Also, we thank the Springer Verlag for publishing this volume and thus helping to position the local activities into an international frame. We hope that this series of publications is contributing to the global promotion of high performance scientific computing.

Stuttgart, Oktober 2007

Wolfgang E. Nagel
Dietmar Kröner
Michael Resch

Contents

Physics

Dr. Roland Speith

Institut für Astronomie und Astrophysik, Abteilung Computational Physics,
Universität Tübingen, Auf der Morgenstelle 10, D-72076 Tübingen, Germany

In the following section, a few projects with general physical background have been selected to give an impression of the applications for physical research currently running at the HLRS. The presented work clearly demonstrates the need for super-computing resources to achieve the respective research objectives, even though the different projects cover a wide range of Physics.

As in the previous years, nearly half of all Physics projects are related to astrophysical research. Exemplary in this field is the work by Marek et al., who report their progress in modelling core collapse supernovae. They investigate the influence of details of the progenitor star on the supernova explosion, in particular non-symmetric effects like stellar rotation. Their results indicate that the neutrino-driven explosion might develop much later after the first shock formation than previously assumed.

Another highlight in the astrophysical field is the project by Brügmann et al., which deals with the simulation of gravitational wave signals generated by the mergers of two orbiting black holes. They managed to perform the largest numerical parameter study to date of black-hole binaries with detailed analysis of the resulting gravitational wave profiles, which are very important as signal templates for the current and future gravitational wave detectors.

Next to Astrophysics, research associated in the broader sense to the field of Quantum Physics dominates the currently running projects. A particular example is the work by Kühn, Steinhauser & Tentyukov regarding theoretical Particle Physics. Their aim is to calculate fundamental parameters of the Standard Model with very high accuracy, which may be helpful to identify deviations from the Standard Model by comparing with high precision experiments.

Finally, as an example for a non-astrophysical, non-Quantum Physics project, Hecht & Harting report on simulations of colloidal particle suspensions in shear flow. They determine the dependence of clustering processes on effects like pH-values and shear rates with a new parallel simulation method.

All the above listed projects once again show that high performance computing has in general become a matter of routine. Also the number of com-

plaints about technical problems is lower compared to the previous year, even though there still have been a few records on code incompatibilities with compilers and unstable runs on particular machines in the reports. Striking is the number of long-term projects which continuously have progressed over years. Roughly three quarters of all Physics projects are directly or indirectly continued from last year. This may demonstrate the on average high excellence of the projects, but it also may indicate a weakness in attracting new research projects and groups to the super-computing centres.

The SuperN-Project: Current Progress in Modelling Core Collapse Supernovae

A. Marek, K. Kifonidis, H.-Th. Janka, and B. Müller

Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, Postfach 1317, D-85741 Garching bei München, Germany
amarek@mpa-garching.mpg.de

Summary. We give an overview of the problems and the current status of our two-dimensional (core collapse) supernova modelling, and discuss the system of equations and the algorithm for its solution that are employed in our code. In particular we report our recent progress, and focus on the ongoing calculations that are performed on the NEC SX-8 at the HLRS Stuttgart. Especially, we will argue that it might be possible that neutrino-driven supernova explosions set in at much later times than previously considered. This, of course, enhances the need of a code that can make efficient use of the multi-node capability of the NEC SX-8 for long-time simulations of the postbounce evolution of collapsing stellar cores.

1 Introduction

A star more massive than about 8 solar masses ends its live in a cataclysmic 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 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 – poorly known – equation of state (EoS) for the proto neutron star? Do softer or stiffer EoSs favour the explosion of a core collapse supernova?
- What is the role of rotation during the explosion? How rapidly do newly formed neutron stars rotate?
- 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 from a future measurement of such signals in case of a Galactic supernova?

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 looses 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 needs to exploit 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 antineutrinos by free nucleons in the post shock layer is thought to reenergize the shock, and lead to 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, 3], 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., [4, 5, 6]), neutrino energy transfer to the postshock 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 work cited above:

(1) Convective instabilities in the neutrino-heated layer between the neutron star and the supernova shock develop to violent convective overturn [7]. 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, and (b) cool matter is able to penetrate closer to the neutron star where it can absorb neutrino energy more efficiently. Both effects allow multi-dimensional models to explode easier than spherically symmetric ones [8, 9, 10].
(2) Recent work [11, 12, 13, 14] has demonstrated that the stalled supernova shock is also subject to a second non-radial low-mode instability, called SASI, which can grow to a dipolar, global deformation of the shock [14, 15].
(3) Convective energy transport inside the nascent neutron star [16, 17, 18, 19] 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 awaits more detailed exploration in 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. All published successful simulations of supernova explosions by the convectivly aided neutrino-heating mechanism in two [8, 9, 20] and three dimensions [21, 22] used such a radical approximation of the neutrino transport.

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 Recent Calculations and the Need for TFlop Simulations

We have recently advanced to a new level of accuracy for supernova simulations by generalising the VERTEX code, a Boltzmann solver for neutrino transport, from spherical symmetry [23] to multi-dimensional applications [24, 25]. The corresponding mathematical model, and in particular our method for tackling the integro-differential transport problem in multi-dimensions, will be summarised in Sect. 3.

Results of a set of simulations with our code in 1D and 2D for progenitor stars with different masses have recently been published by [25, 26], and with respect to the expected gravitational-wave signals from rotating and convective supernova cores by [27]. The recent progress in supernova modelling was summarised and set in perspective in a conference article by [24].

Our collection of simulations has helped us to identify a number of effects which have brought our two-dimensional models close to the threshold of explosion. This makes us optimistic that the solution of the long-standing problem of how massive stars explode may be in reach. In particular, we have recognised the following aspects as advantageous:

- The details of the stellar progenitor (i.e. the mass of the iron core and its radius–density relation) have substantial influence on the supernova evolution. Especially, we found explosions of stellar models with low-mass (i.e. small) iron cores [26, 28], whereas more massive stars resist the explosion more persistent [25]. Thus detailed studies with different progenitor models are necessary.
- Stellar rotation, even at a moderate level, supports the expansion of the stalled shock by centrifugal forces and instigates overturn motion in the neutrino-heated postshock matter by meridional circulation flows in addition to convective instabilities.

All these effects are potentially important, and some (or even all of them) may represent crucial ingredients for a successful supernova simulation. So far no multi-dimensional calculations have been performed, in which two or more of these items have been taken into account simultaneously, and thus their mutual interaction awaits to be investigated. It should also be kept in mind that our knowledge of supernova microphysics, and especially the EoS of neutron star matter, is still incomplete, which implies major uncertainties for supernova modelling. Unfortunately, the impact of different descriptions for this input physics has so far not been satisfactorily explored with respect to the neutrino-heating mechanism and the long-time behaviour of the supernova shock, in particular in multi-dimensional models. However, first

multi-dimensional simulations of core collapse supernovae with different nuclear EoSs [19, 29] show a strong dependence of the supernova evolution on the EoS.

From this it is clear that rather extensive parameter studies using multi-dimensional simulations are required to identify the physical processes which are essential for the explosion. Since on a dedicated machine performing at a sustained speed of about 30 GFlops already a single 2D simulation has a turn-around time of more than half a year, these parameter studies are not possible without TFlop simulations.

3 The Mathematical Model

The non-linear system of partial differential equations which is solved in our code consists of the following components:

- The Euler equations of hydrodynamics, supplemented by advection equations for the electron fraction and the chemical composition of the fluid, and formulated in spherical coordinates;
- the Poisson equation for calculating the gravitational source terms which enter the Euler equations, including corrections for general relativistic effects;
- the Boltzmann transport equation which determines the (non-equilibrium) distribution function of the neutrinos;
- the emission, absorption, and scattering rates of neutrinos, which are required for the solution of the Boltzmann equation;
- 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. For a more complete description of the entire code we refer the reader to [25], and the references therein.

3.1 "Ray-by-Ray Plus" Variable Eddington Factor Solution of 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 seven-dimensional functions, 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

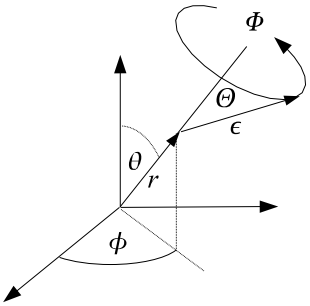

Fig. 1. Illustration of the phase space coordinates (see the main text)

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 dynamics of the flow. It makes therefore 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

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

where $\mathrm{d}\Omega = \sin\Theta\, \mathrm{d}\Theta\, \mathrm{d}\Phi$, $\boldsymbol{n} = (\cos\Theta, \sin\Theta\cos\Phi, \sin\Theta\sin\Phi)$, 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 resort to assumptions on its symmetry. At this point, one usually employs azimuthal symmetry for the stellar matter distribution, i.e. any dependence on the azimuth angle ϕ is ignored, which implies that the hydrodynamics of the problem can be treated in two dimensions. It also implies $I(r, \vartheta, \epsilon, \Theta, \Phi) = I(r, \vartheta, \epsilon, \Theta, -\Phi)$. If, in addition, it is assumed that I is even independent of Φ, then each of the angular moments of I becomes a *scalar*, which depends on two spatial dimensions, and one dimension in momentum space: $J, H, K, L = J, H, K, L(r, \vartheta, \epsilon, t)$. Thus we have reduced the problem to three dimensions in total.

The System of Equations

With the aforementioned assumptions it can be shown [25], that in order to compute the source terms for the energy and electron fraction of the fluid, the following two transport equations need to be solved:

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

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

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 $\beta_r = v_r/c$ and $\beta_\vartheta = v_\vartheta/c$, where v_r and v_ϑ are the components of the hydrodynamic velocity, and c is the speed of light. The functional dependences $\beta_r = \beta_r(r,\vartheta,t)$, $J = J(r,\vartheta,\epsilon,t)$, 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 is indeed determined from a separate simplified ("model") Boltzmann equation.

A finite volume discretisation of (2–3) is sufficient to guarantee exact conservation of the total neutrino energy. However, and as described in detail in [23], it is not sufficient to guarantee also exact conservation of the neutrino number. To achieve this, we discretise and solve a set of two additional equations. With $\mathcal{J} = J/\epsilon$, $\mathcal{H} = H/\epsilon$, $\mathcal{K} = K/\epsilon$, and $\mathcal{L} = L/\epsilon$, this set of equations reads

$$\left(\frac{1}{c}\frac{\partial}{\partial t}+\beta_r\frac{\partial}{\partial r}+\frac{\boldsymbol{\beta_\vartheta}}{\boldsymbol{r}}\frac{\boldsymbol{\partial}}{\boldsymbol{\partial\vartheta}}\right)\mathcal{J}+\mathcal{J}\left(\frac{1}{r^2}\frac{\partial(r^2\beta_r)}{\partial r}+\frac{\mathbf{1}}{\boldsymbol{r\sin\vartheta}}\frac{\boldsymbol{\partial(\sin\vartheta\beta_\vartheta)}}{\boldsymbol{\partial\vartheta}}\right)$$
$$+\frac{1}{r^2}\frac{\partial(r^2\mathcal{H})}{\partial r}+\frac{\beta_r}{c}\frac{\partial\mathcal{H}}{\partial t}-\frac{\partial}{\partial\epsilon}\left\{\frac{\epsilon}{c}\frac{\partial\beta_r}{\partial t}\mathcal{H}\right\}-\frac{\partial}{\partial\epsilon}\left\{\epsilon\mathcal{J}\left(\frac{\beta_r}{r}+\frac{\mathbf{1}}{\boldsymbol{2r\sin\vartheta}}\frac{\boldsymbol{\partial(\sin\vartheta\beta_\vartheta)}}{\boldsymbol{\partial\vartheta}}\right)\right\}$$
$$-\frac{\partial}{\partial\epsilon}\left\{\epsilon\mathcal{K}\left(\frac{\partial\beta_r}{\partial r}-\frac{\beta_r}{r}-\frac{\mathbf{1}}{\boldsymbol{2r\sin\vartheta}}\frac{\boldsymbol{\partial(\sin\vartheta\beta_\vartheta)}}{\boldsymbol{\partial\vartheta}}\right)\right\}+\frac{1}{c}\frac{\partial\beta_r}{\partial t}\mathcal{H}=\mathcal{C}^{(0)}, \quad (4)$$

$$\left(\frac{1}{c}\frac{\partial}{\partial t}+\beta_r\frac{\partial}{\partial r}+\frac{\boldsymbol{\beta_\vartheta}}{\boldsymbol{r}}\frac{\boldsymbol{\partial}}{\boldsymbol{\partial\vartheta}}\right)\mathcal{H}+\mathcal{H}\left(\frac{1}{r^2}\frac{\partial(r^2\beta_r)}{\partial r}+\frac{\mathbf{1}}{\boldsymbol{r\sin\vartheta}}\frac{\boldsymbol{\partial(\sin\vartheta\beta_\vartheta)}}{\boldsymbol{\partial\vartheta}}\right)$$
$$+\frac{\partial\mathcal{K}}{\partial r}+\frac{3\mathcal{K}-\mathcal{J}}{r}+\mathcal{H}\left(\frac{\partial\beta_r}{\partial r}\right)+\frac{\beta_r}{c}\frac{\partial\mathcal{K}}{\partial t}-\frac{\partial}{\partial\epsilon}\left\{\frac{\epsilon}{c}\frac{\partial\beta_r}{\partial t}\mathcal{K}\right\}$$
$$-\frac{\partial}{\partial\epsilon}\left\{\epsilon\mathcal{L}\left(\frac{\partial\beta_r}{\partial r}-\frac{\beta_r}{r}-\frac{\mathbf{1}}{\boldsymbol{2r\sin\vartheta}}\frac{\boldsymbol{\partial(\sin\vartheta\beta_\vartheta)}}{\boldsymbol{\partial\vartheta}}\right)\right\}$$
$$-\frac{\partial}{\partial\epsilon}\left\{\epsilon\mathcal{H}\left(\frac{\beta_r}{r}+\frac{\mathbf{1}}{\boldsymbol{2r\sin\vartheta}}\frac{\boldsymbol{\partial(\sin\vartheta\beta_\vartheta)}}{\boldsymbol{\partial\vartheta}}\right)\right\}-\mathcal{L}\left(\frac{\partial\beta_r}{\partial r}-\frac{\beta_r}{r}-\frac{\mathbf{1}}{\boldsymbol{2r\sin\vartheta}}\frac{\boldsymbol{\partial(\sin\vartheta\beta_\vartheta)}}{\boldsymbol{\partial\vartheta}}\right)$$
$$-\mathcal{H}\left(\frac{\beta_r}{r}+\frac{\mathbf{1}}{\boldsymbol{2r\sin\vartheta}}\frac{\boldsymbol{\partial(\sin\vartheta\beta_\vartheta)}}{\boldsymbol{\partial\vartheta}}\right)+\frac{1}{c}\frac{\partial\beta_r}{\partial t}\mathcal{J}=\mathcal{C}^{(1)}. \quad (5)$$

The moment equations (2–5) are very similar to the $\mathcal{O}(v/c)$ equations in spherical symmetry which were solved in the 1D simulations of [23] (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 VERTEX, 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

$$\frac{\mathrm{d}e}{\mathrm{d}t}=-\frac{4\pi}{\rho}\int_0^\infty \mathrm{d}\epsilon \sum_{\nu\in(\nu_\mathrm{e},\bar{\nu}_\mathrm{e},\dots)} C_\nu^{(0)}(\epsilon), \quad (6)$$

$$\frac{\mathrm{d}Y_\mathrm{e}}{\mathrm{d}t}=-\frac{4\pi\, m_\mathrm{B}}{\rho}\int_0^\infty \mathrm{d}\epsilon \left(\mathcal{C}_{\nu_\mathrm{e}}^{(0)}(\epsilon)-\mathcal{C}_{\bar{\nu}_\mathrm{e}}^{(0)}(\epsilon)\right) \quad (7)$$

(for the momentum source terms due to neutrinos see [25]). Here m_B is the baryon mass, and the sum in (6) runs over all neutrino types. The full system consisting of (2–7) is stiff, and thus requires an appropriate discretisation scheme for its stable solution.

Method of Solution

In order to discretise (2–7), the spatial domain $[0, r_\mathrm{max}] \times [\vartheta_\mathrm{min}, \vartheta_\mathrm{max}]$ is covered by N_r radial, and N_ϑ angular zones, where $\vartheta_\mathrm{min} = 0$ and $\vartheta_\mathrm{max} = \pi$ correspond to the north and south poles, respectively, of the spherical grid. (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 two operator-split steps corresponding to a lateral and a radial sweep.

In the first step, we treat the boldface terms in the respectively first lines of (2–5), which describe the lateral 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 equation

$$\frac{1}{c}\frac{\partial \Xi}{\partial t} + \frac{1}{r\sin\vartheta}\frac{\partial(\sin\vartheta\,\beta_\vartheta\,\Xi)}{\partial\vartheta} = 0\,, \tag{8}$$

where Ξ represents one of the moments J, H, $\mathcal{J}$, or $\mathcal{H}$. 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 second 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 discretisation 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, H) or $(\mathcal{J}, \mathcal{H})$. They only depend on the hydrodynamic velocity v_ϑ, which is 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 (2–5) stem from the Boltzmann equation, which is an integro-differential equation and couples all the energy bins
- the discretisation 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 angle ϑ and for each type of neutrino, i.e. for constant ϑ, N_ν two-dimensional problems need to be solved.

The discretisation itself is done using a second order accurate scheme with backward differencing in time according to [23]. 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.

4 Recent Results and Ongoing Work

We make use of the computer resources available to us at the HLRS to address some of the important questions in SN theory (see Sect. 1) with 2D-simulations. At the HLRS, we typically run our code on one node of the NEC SX-8 (8 processors, OpenMP-parallised) of the SX-8 with 98.3% of vector operations and 22000 MFLOPS per second.

In the following we present some of our results from these simulations that are currently performed at the HLRS. For the neutrino interaction rates we use the full set as described in [30] and general relativistic effects are taken into account according to [31].

4.1 The Importance of Hydrodynamic Instabilities Inside the Supernova Core

Buras et al. [26] recently reported two simulations of an 11.2 $M_{\odot}$ stellar progenitor model with two different setups: on the one hand they calculated a 90° wedge centred around the equatorial plane with periodic boundary conditions (from here on called model s11.2-wedge). On the other hand a full 180° model (from north to south pole) with reflecting boundary conditions was calculated (from here on this model is called s11.2-180). Interestingly, the latter model showed an explosion whereas the former one failed to explode.

In order to investigate whether the different wedge size or the different boundary conditions caused this qualitative difference of the simulations, we have performed simulations for two additional models: the first one, model s11.2-90, employed a 90° wedge from the pole to the equator with reflecting boundary conditions. The second model, s11.2-wedge-refl, was chosen with a setup as model s11.2-wedge, however, reflecting boundary conditions were used. The basic setup of all four models is sketched in Fig. 2.

We found that both models containing the polar axis developed explosions and that the wedge models did not explode *independent* of the boundary

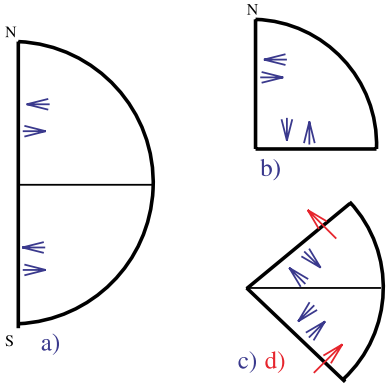

Fig. 2. Schematic overview over the setups chosen for model s11.2-180 (**a**), s11.2-90 (**b**), s11.2-wedge-refl (**c**), and s11.2-wedge (**d**). *Blue arrows* indicate the use of reflecting boundary conditions, whereas *red arrows* indicate the use of periodic boundary conditions. The sketches are oriented such that the equatorial plane is represented by the *horizontal line*

condition, although all models show a similar strength of neutrino heating. A more detailed analysis [19] showed that in the exploding models a strong non-radial hydrodynamical instability of the accretion shock, the so-called SASI [13], developed, which was suppressed in the non-exploding (wedge) models. A projection of the shock-positions onto spherical harmonics $Y_{l,m}$,

$$r_{\rm shock}(\theta, t) = \sum_{l=0}^{\infty} a_{l,0}(t) Y_{l,0}(\theta) \ , \tag{9}$$

where m is set to zero due to the adopted axial symmetry of our 2D calculations, clearly shows the presence of this time and angle dependent shock instability, see Fig. 3.

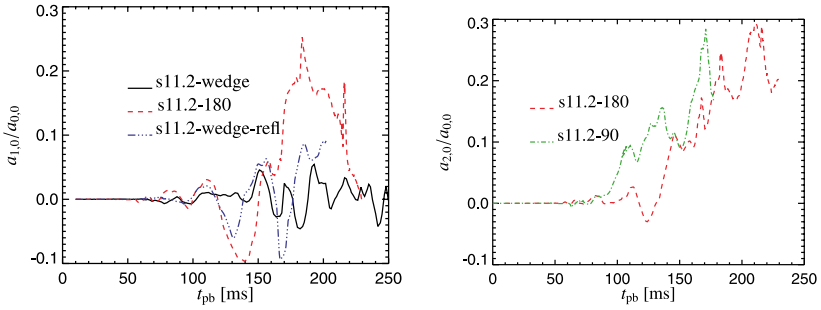

Fig. 3. Time-dependent coefficients of the expansion of the shock position into spherical harmonics (cf. (9)) Note that the the coefficients are normalised to the amplitude of the $l = 0$ mode. **a**: The coefficients corresponding to the $l = 1$ mode. **b**: The coefficients corresponding to the $l = 2$ mode

From this, together with (semi) analytical arguments given in [12, 13] (see also Sect. 2.1), we conclude that the SASI can play a crucial role for the development of an explosion. Indeed up to now all multi-dimensional explosion models of the Garching supernova group showed such kind if instabilities gaining strength before the onset of the explosion (see also paragraph 4.2).

4.2 Effects of Rotation

As we have already described in our last report [29], we have started a full 180° simulation of a 15 $M_\odot$ progenitor star including rotation. This still ongoing simulation, model L&S-rot, – it was partly performed at the SX-8 of the HLRS – is now the longest ($t_{\rm pb} > 600$ ms) multidimensional Boltzmann neutrino transport simulation worldwide. As we argued in our last report, the reason for pushing this simulation to such late times is that rotation and angular momentum become more and more important at later times as matter has

fallen from larger radii to the shock position. During all the simulated supernova evolution we find the presence of the already mentioned SASI low-mode deformation of the accretion layer and shock (see previous Sect. 4.1), and at a time of roughly 500 ms after the shock formation we observe the onset of an explosion (see Fig. 4). However, it is yet not clear whether the presence of rotation is crucial for the explosion of this 15 $M_\odot$ model, or whether this

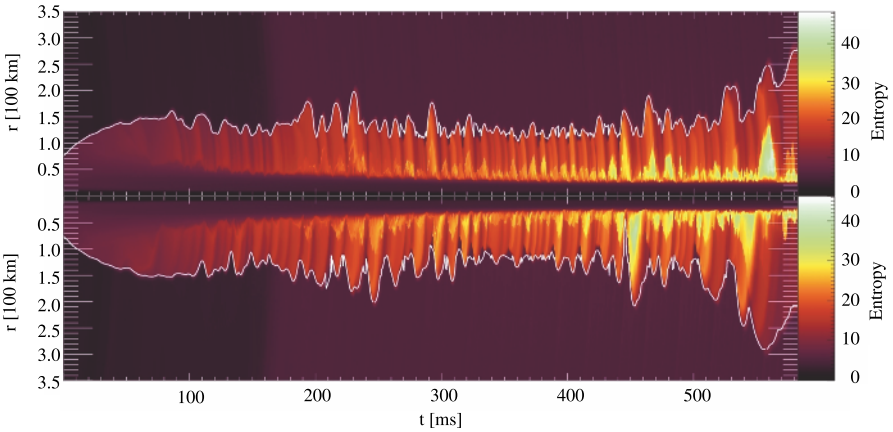

Fig. 4. The shock position (*solid white line*) at the north pole (*upper panel*) and south pole (*lower panel*) of the rotating 15 $M_\odot$ model as function of postbounce time. Colour coded is the entropy of the stellar fluid

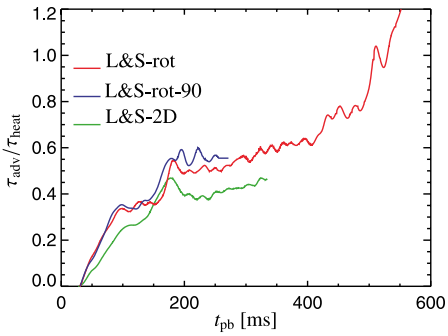

Fig. 5. The ratio of the advection timescale to the heating timescale for the rotating model L&S-rot and the non-rotating model L&S-2D. Also shown is model L&S-rot-90 which is identical to model L&S-rot except for the computational domain that does not extend from pole to pole but from the north pole to the equator (see also Sect. 4.1). The advection timescale is the characteristic timescale that matter stays inside the heating region before it is advected to the proto-neutron star. The heating timescale is the typical timescale that matter needs to be exposed to neutrino heating for observing enough energy to become gravitationally unbound

model would also explode without rotation. Since the comparison of the rotating and a corresponding non-rotating model reveals qualitatively the same behaviour, see e.g. Fig. 5, it is absolutely necessary to evolve both models to a time of more than 500 ms after the shock formation in order to answer this question.

Although the importance of rotation in this particular model is not yet fully understood, this calculation nevertheless shows that the convectivly or SASI supported neutrino-heating mechanism is viable to produce explosions for massive progenitors ($M > 10\,\mathrm{M}_\odot$) much later than previously considered.

5 Conclusions and Outlook

We continued to simulate well resolved 2D models of core collapse supernovae with detailed neutrino transport at the HLRS. We found that non-radial hydrodynamic instabilities support the onset of supernova explosions and for a $15\,\mathrm{M}_\odot$ progenitor model we obtained a supernova explosion at a time of roughly 600 ms after the shock formation. These results indicate that supernova explosions might develop at much later times than previously thought, which has to be investigated in more detail in the future. Especially the non-rotating model has to be evolved to a time of roughly 700 ms after the shock formation, in order to clarify the importance of rotation in the exploding model. All in all, this means longer evolution times per model and thus emphasises the need for very efficient numerical tools. With support by the staff of the HLRS, our code is currently optimised for multi-node usage on the NEC SX-8 which will allow us to evolve our supernova models faster to longer evolution times than it is possible at the moment. This will be crucial for investigating the importance of the stellar rotation as well as the possibility of supernova explosions at late times ($t > 0.5\,\mathrm{s}$).

Acknowledgements

Support from the SFB 375 “Astroparticle Physics”, SFB/TR7 “Gravitationswellenastronomie”, and SFB/TR27 “Neutrinos and Beyond” of the Deutsche Forschungsgemeinschaft, and computer time at the HLRS and the Rechenzentrum Garching are acknowledged. We also thank M. Galle and R. Fischer for performing the benchmarks on the NEC machines. We thank especially K. Benkert for further optimising our code for the SX-8 architecture. The simulations were performed on the national super computer NEXC SX-8 at the High Performance Computing Center Stuttgart (HLRS) under the grant number SuperN/12758.

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. Mezzacappa, A., Liebendörfer, M., Messer, O.E., Hix, W.R., Thielemann, F., Bruenn, S.W.: Simulation of the Spherically Symmetric Stellar Core Collapse, Bounce, and Postbounce Evolution of a Star of 13 Solar Masses with Boltzmann Neutrino Transport, and Its Implications for the Supernova Mechanism. Phys. Rev. Letters **86** (2001) 1935–1938
3. 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–+
4. Bethe, H.A.: Supernova mechanisms. Reviews of Modern Physics **62** (1990) 801–866
5. Burrows, A., Goshy, J.: A Theory of Supernova Explosions. Astrophys. J. **416** (1993) L75
6. Janka, H.T.: Conditions for shock revival by neutrino heating in core-collapse supernovae. Astron. Astrophys. **368** (2001) 527–560
7. Herant, M., Benz, W., Colgate, S.: Postcollapse hydrodynamics of SN 1987A – Two-dimensional simulations of the early evolution. Astrophys. J. **395** (1992) 642–653
8. 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
9. Burrows, A., Hayes, J., Fryxell, B.A.: On the nature of core-collapse supernova explosions. Astrophys. J. **450** (1995) 830
10. Janka, H.T., Müller, E.: Neutrino heating, convection, and the mechanism of Type-II supernova explosions. Astron. Astrophys. **306** (1996) 167–+
11. Thompson, C.: Accretional Heating of Asymmetric Supernova Cores. Astrophys. J. **534** (2000) 915–933
12. Foglizzo, T.: Non-radial instabilities of isothermal Bondi accretion with a shock: Vortical-acoustic cycle vs. post-shock acceleration. Astron. Astrophys. **392** (2002) 353–368
13. 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
14. 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–+
15. Scheck, L.: Multidimensional simulations of core collapse supernovae. PhD thesis, Technische Universität München (2006)
16. Keil, W., Janka, H.T., Mueller, 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. Marek, A.: Multi-dimensional simulations of core collapse supernovae with different equations of state for hot proto-neutron stars. PhD thesis, Technische Universität München (2007)
20. Fryer, C.L., Heger, A.: Core-Collapse Simulations of Rotating Stars. Astrophys. J. **541** (2000) 1033–1050
21. Fryer, C.L., Warren, M.S.: Modeling Core-Collapse Supernovae in Three Dimensions. Astrophys. J. **574** (2002) L65–L68
22. Fryer, C.L., Warren, M.S.: The Collapse of Rotating Massive Stars in Three Dimensions. Astrophys. J. **601** (2004) 391–404
23. Rampp, M., Janka, H.T.: Radiation hydrodynamics with neutrinos. Variable Eddington factor method for core-collapse supernova simulations. Astron. Astrophys. **396** (2002) 361–392
24. Janka, H.T., Buras, R., Kifonidis, K., Marek, A., Rampp, M.: Core-Collapse Supernovae at the Threshold. In Marcaide, J.M., Weiler, K.W., eds.: Supernovae, Procs. of the IAU Coll. 192, Berlin, Springer (2004)
25. 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 $15\,M_\odot$ star. Astron. Astrophys. **447** (2006) 1049–1092
26. Buras, R., Janka, H.T., Rampp, M., Kifonidis, K.: Two–dimensional hydrodynamic core–collapse supernova simulations with spectral neutrino transport. II. Models for different progenitor stars. Astron. Astrophys. **457** (2006) 281–308
27. Müller, E., Rampp, M., Buras, R., Janka, H.T., Shoemaker, D.H.: Toward Gravitational Wave Signals from Realistic Core-Collapse Supernova Models. Astrophys. J. **603** (2004) 221–230
28. Kitaura, F.S., Janka, H.T., Hillebrandt, W.: Explosions of O–Ne–Mg cores, the Crab supernova, and subluminous type II–P supernovae. Astron. Astrophys. **450** (2006) 345–350
29. Marek, A., Kifonidis, K., Janka, H.T., Müller, B.: The supern-project: Understanding core collapse supernovae. In Nagel, W.E., Jäger, W., Resch, M., eds.: High Performance Computing in Science and Engineering 06, Berlin, Springer (2006)
30. Marek, A., Janka, H.T., Buras, R., Liebendörfer, M., Rampp, M.: On ion-ion correlation effects during stellar core collapse. Astron. Astrophys. **443** (2005) 201–210
31. Marek, A., Dimmelmeier, H., Janka, H.T., Müller, E., Buras, R.: Exploring the relativistic regime with Newtonian hydrodynamics: an improved effective gravitational potential for supernova simulations. Astron. Astrophys. **445** (2006) 273–289

Toward Conquering the Parameter Space of Gravitational Wave Signals from Black Hole Coalescence

Bernd Brügmann[1], Jose Gonzalez[1], Mark Hannam[1], Sascha Husa[1], Ulrich Sperhake[1], and Iris Christadler[2]

[1] Theoretisch Physikalisches Institut, Universität Jena, Max-Wien-Platz 1, 07743 Jena, Germany
sascha.husa@uni-jena.de
[2] Leibniz Rechenzentrum der Bayerischen Akademie der Wissenschaften, Boltzmannstraße 1, 85748 Garching, Germany

1 Introduction

We present a status report on our second year of work on the numerical simulation of gravitational wave signals from coalescing black holes at HLRS. We report in particular on a significant number of important results we have obtained, and on progress with optimizing our algorithms and code.

Gravitational waves and their sources are described by the Einstein equations, which underly the theory of general relativity (see e.g., our previous report [23], or the overview article [31]). A large international effort is currently underway to detect gravitational waves and start a new field of astrophysical research, gravitational wave astronomy. Within Germany, this research is conducted within the SFB/TR 7 "Gravitational Wave Astronomy".

In our work we solve the Einstein equations numerically, mainly with finite-difference methods. Essentially, the solution process has two parts: first, appropriate initial data have to be constructed by solving elliptic constraint equations, which in some sense generalize the divergence constraints of electromagnetism. Second, these data are evolved by solving a system of hyperbolic equations. The hyperbolic character of the evolution equations expresses the fact that in general relativity physical information propagates at the speed of light.

The particular type of gravitational wave source that we study is the inspiral and coalescence of black-hole binaries. The efficiency of gravitational wave sources is related to their compactness: A weak field calculation yields the loss of energy of a system as

$$\frac{dE}{dt} = -\frac{G}{5c^5} \sum_{i,j} \left(\frac{d^3 Q_{ij}}{dt^3} \right)^2 \tag{1}$$

where $Q_{ij} = \int \varrho(x_i x_j - \frac{1}{3}\delta_{ij} r^2) d^3x$ is the mass quadrupole moment. The radiation power thus scales with the sixth power of the frequency of the system. Due to the weakness of gravity, expressed in the factor $\frac{G}{5c^5}$, only systems of astrophysical dimensions – large masses moving at a significant fraction of the speed of light – generate significant amounts of gravitational radiation. Since black holes are the most compact objects we can consider, black-hole binaries are a particularly efficient and important source for current and future gravitational wave detectors.

Gravitational waves have so far been detected only indirectly by comparing the tightening of the orbit of the Hulse-Taylor binary pulsar with the theoretical prediction from energy loss due to gravitational wave emission [41, 54]. However, a growing network of gravitational-wave detectors such as LIGO [1, 43], GEO [29, 33] and VIRGO [56] is currently taking science data and direct detection is forseen for the next few years. In order to actually extract information about the sources from observations, accurate signal templates are needed for various types of sources.

Producing templates for gravitational-wave data analysis will require large parameter studies, and correspondingly large computational resources: The eventual goal of our simulations is to map the parameter space of gravitational wave signals from black hole coalescence. Our simulations typically start with initial data that correspond to the astrophysically most relevant case of quasi-circular inspiral, which essentially means that the eccentricity is very small. The parameter space is then essentially given by the mass ratio and individual spins, as well as the initial orientation of the spins. The latter determines in particular the spin orientation at merger time, which has a strong influence on the gravitational wave signal. In order to produce "complete" waveforms, which contain large numbers of gravitational wave cycles from the inspiral phase, as well as the merger and ringdown phases, it is necessary to start the numerical simulations in the regime where Post-Newtonian analytical calculations are valid. These describe very accurately the waveforms of the early inspiral process, but break down for small separations of the black holes. This "matching" of analytical and numerical results requires large initial black-hole separations and large integration times. Improving current estimates for the validity of the Post-Newtonian approximation is an integral part of the research, and will eventually determine the computational cost.

Crucial technical problems in the field at this point relate to the efficiency of the simulations, and to the establishment of a "data analysis pipeline", connecting analytical calculations of the early inspiral phase with numerical simulations and gravitational-wave searches in actual detector data.

Since our last report, we have obtained a large number of important results:

- Through careful analysis of simulations of a single black hole, we described for the first time the geometrical behavior of numerical solutions in the now-standard "moving puncture" method of evolving black-hole spacetimes [38, 39].
- We have presented the most detailed study of the two proposed versions of the moving-puncture method as applied to equal-mass binaries [21].
- All moving-puncture simulations to date have used a particular type of initial data, Bowen-York puncture data [19, 20]. These data make the physically unrealistic assumption of conformal flatness. We have made significant progress in removing this assumption for binaries made up of spinning black holes [37].
- An independent code, LEAN, has been developed, and used to perform a comparison of head-on collision simulations using puncture and excision initial data [53].
- Astrophysicists have for many years been waiting for numerical simulations to accurately calculate the recoil when two black holes of unequal mass collide. In the largest parameter study to date of black-hole binaries requiring approximately 150000 CPU hours, we have calculated the maximum recoil of unequal-mass nonspinning binaries to be 175.2 ± 11 km/s [35], and performed a detailed analysis of the gravitational waves emitted during unequal-mass binary mergers [17].
- Far more spectacularly, we have found that extremely large recoils are possible for spinning black holes, and estimated the kick from one configuration as 2500 km/s, large enough to eject the remnant black hole even from a giant elliptical galaxy [34].
- We have begun work on establishing a data-analysis pipeline to provide our numerical waveforms to the gravitational-wave data-analysis community [2, 3]. This work may increase the chances of gravitational-wave detection by at least a factor of ten.

Further research in all of these areas is still underway; there are a wealth of areas of explore. In this paper we will provide further details of these results, and present performance tuning results and a summary of code optimization over the last twelve months.

2 Status of the Field

The numerical solution of the full Einstein equations represents a very complex problem, and for two black holes the spacetime singularities that are encountered in the interior of black holes pose an additional challenge. In order to obtain accurate results both the use of mesh refinement techniques *and* a good choice of coordinate gauge are essential. Together with the complicated structure of the equations – a typical code has between ten and several dozen evolution variables, and, when expanded, the right hand sides of the

evolution equations have thousands of terms – this yields a computationally very complex and mathematically very subtle problem.

For a long time, typical runs had been severely limited by the achievable evolution time before the simulations become too inaccurate or before the computer code became unstable, and there were serious doubts whether numerical relativity techniques could produce gravitational-wave templates, at least in the near future. This picture has drastically changed ever since the first simulations of a complete black hole orbit were obtained in early 2004 [24]. In spring of 2005 Pretorius [49] presented the first simulation lasting for several orbits, using adaptive mesh refinement, second-order finite differencing, a sophisticated method to excise the singular interior of the black hole from the grid, and an implicit evolution algorithm.

An alternative to the "excision" method of treating black holes is to "fill" the black hole with a topological defect in the form of an interior spacelike asymptotic end, the "puncture" [5, 20], and to freeze the evolution of the asymptotic region through a judicious choice of coordinate gauge [6, 7, 10, 22]. The latter approach, combined with a setup where the topological defect is allowed to move across the grid ("moving puncture" approach [11, 28]) has lead to a giant leap forward in the field [12, 13, 26, 27, 39, 40, 53], taking the first orbit simulations of black holes [24, 49] to more than ten orbits and allowing accurate wave extraction. It is this approach of moving punctures that we and most groups follow. There are now ten groups capable of performing black hole binary evolutions, with Jena being one of five groups leading the field.

Recent work in the field has shifted focus from methodological work to studies of astrophysical relevance. A particular focus of the last few months has been the on the so-called recoil or rocket effect due to "beamed" emission of gravitational radiation [15, 18, 48]. By momentum conservation, radiation of energy in a preferred direction corresponds to a loss of linear momentum and the black hole that results from the merger thus recoils from the center-of-mass frame with speeds of up to several thousand km/s. The velocity of this "kick" depends on the configuration of the system (e.g. the mass ratios and spins) and details of the merger dynamics, but not on the total mass (velocity is dimensionless in geometric units).

From an astrophysical point of view, the recoil effect is particularly interesting for massive black holes with masses $> 10^5\, M_\odot$, which exist at the center of many galaxies and have a substantial impact on the structure and formation of their host galaxies, as is demonstrated by the correlation of the black-hole mass with the bulge mass, luminosity and velocity dispersion [32, 45, 46, 47]. The largest recoil effects have so far been found [34] for a particularly simple configuration: equal mass black holes with (initially) anti-aligned spins in the orbital plane. Such large kicks are on the order of 1% of the speed of light, and larger than the escape velocity of about 2000 km/s of giant elliptical galaxies.

In order to predict accurate gravitational wave signals, it is necessary to perform simulations over many orbits and connect such numerical simula-

tions to post-Newtonian analytical calculations. It turns out that in such long simulations phase errors of the orbits are very hard to control, and accurate simulations with second or fourth order accurate methods are prohibitively expensive. As an example for long evolutions we show results from three simulations with initial coordinate separations $D = 10, 12, 14$, corresponding to somewhat more than 5, 9 and 14 orbits. Orbital tracks are shown in Fig. 1, and gravitational wave signals in Fig. 2. The signals have been shifted in time to reach the maximum at $t = 0$, which in particular compensates for the large numerical errors in phase. In order to overcome phase inaccuracies in long

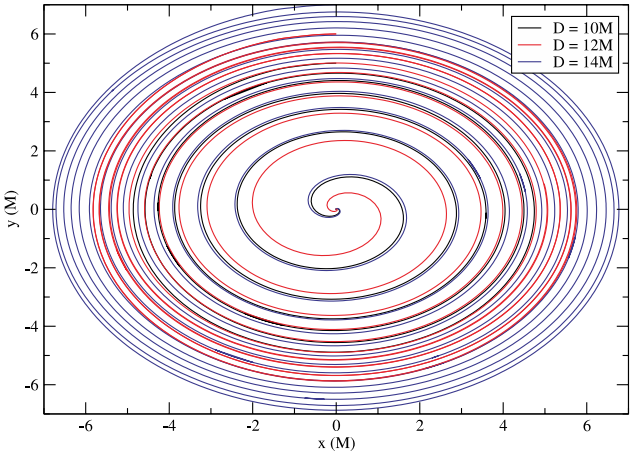

Fig. 1. Coordinate tracks of the puncture locations for inspirals with initial coordinate separations $D = 10, 12, 14$, corresponding to a bit more than 5, 9 and 14 orbits

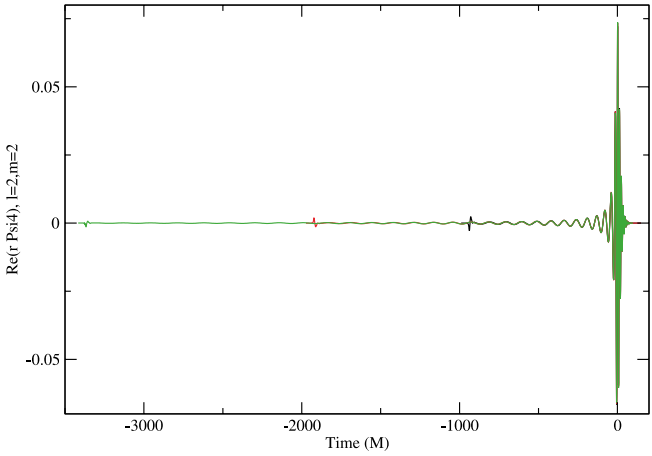

Fig. 2. Gravitational wave signals from the simulations shown in Fig. 1, shifted in time to reach the maximum at $t = 0$

evolutions, spectral methods have been suggested and significant progress has been made by the Caltech-Cornell group [50]. In our group we have recently obtained excellent results with 6th order accurate finite differencing methods as described below. We are currently in the process of fine-tuning our grid setup for 6th order evolutions for different types of initial data (adjustments are required for black holes with spins or unequal mass evolutions).

3 Description of the Method

3.1 Continuum Problem

Writing the Einstein equations as an initial-value problem yields a coupled system of second differential order elliptic constraints and hyperbolic evolution equations that preserve the constraints. In the free evolution approach, which is most common in the field, the constraints are only solved initially, and later only the hyperbolic equations are used to construct the solution. There is a large freedom in writing the Einstein equations as a system of partial differential equations, and much research has gone into finding optimal choices. We currently employ the most popular choice, the so-called BSSN system [4, 7, 14, 36, 52].

3.2 Numerical Methods for the Evolution Problem

Discretization in space is performed with standard fourth-order accurate stencils, although we are currently working on a transition to making sixth order accurate methods our standard choice. Symmetric stencils are used with the exception of the advection terms associated with the shift vector, where we use lop-sided upwind stencils, see e.g. [57] for the fourth order case. For the sixth order case we find that several choices for the advection term stencils yield stable evolutions, but the lop-sided upwind stencil which is closest to the symmetric case yields by far the best accuracy, i.e. we use

$$f'(x) = \frac{2f(x-2h) - 24f(x-h) - 35f(x) + 80f(x+h)}{60h} + \frac{-30f(x+2h) + 8f(x+3h) - f(x+4h)}{60h} - \frac{1}{105}\frac{d^7 f(x)}{dx^7}h^6.$$

Alternative asymmetric choices would be

$$f'(x) = \frac{-10f(x-h) - 77f(x) + 150f(x+h) - 100f(x+2h)}{60h} + \frac{50f(x+3h) - 15f(x+4h) + 2f(x+5h)}{60h} + \frac{1}{42}\frac{d^7 f(x)}{dx^7}h^6,$$

$$f'(x) = \frac{-147f(x) + 360f(x+h) - 450f(x+2h) + 400f(x+3h)}{60h} + \frac{-225f(x+4h) + 72f(x+5h) - 10f(x+6h)}{60h} - \frac{1}{7}\frac{d^7 f(x)}{dx^7}h^6$$

for the stencils that deviate more from the symmetric choice. We can see that the first choice has the smallest leading error term. The symmetric stencil has an even smaller error term $\frac{1}{140}\frac{d^7 f(x)}{dx^7}h^6$, but does not show equally robust results, as is common for solving advection equations. For non-advection derivative terms we again use the standard symmetric stencil. Time integration is performed by standard Runge-Kutta type methods, in particular 3rd and 4th order Runge-Kutta and second-order accurate three-step iterative Crank-Nicholson integrators as described in [25], where Courant limits and stability properties are discussed for the types of equations used here.

Mesh refinement techniques are utilized to resolve the different scales of the problem, and to follow the motion of the black holes. The relevant spatial scales of the binary black-hole problem are the scales of the holes, their orbital motion, the typical wave lengths of the ring-down of the individual and merged black holes, the typical wavelength of the merger waveform and the asymptotic falloff of the fields. All of these scales can be estimated from the initial data and vary relatively slowly with time. It is thus very efficient to essentially use a fixed mesh refinement strategy, with inner level refinement boxes following the motion of the black holes. Typically we use about 10 refinement levels (refining the grid spacing by factors of 2), roughly half of which follow the movement of the black holes. The numerical algorithm is a modification of the standard Berger-Oliger algorithm [16] as described in [21].

3.3 Numerical Methods for Solving the Constraints

Since the Einstein equations constitute a constrained system, initial data for the evolution equations has to satisfy the constraints. The binary black hole initial data used in our simulations is so-called puncture initial data, which requires the solution of one scalar non-linear elliptic equation [20]. In the BAM code such data can be obtained with a parallel geometric multigrid solver, which however is using second order finite differencing.

During the last funding period we have implemented an interface to the special purpose pseudo-spectral collocation code described in [9] for the BAM code. Such data was already available to LEAN as a Cactus module.

In BAM, the AMR data is obtained by barycentric interpolation, typically with 8th order polynomials for both the 4th and 6th order finite differencing methods. The efficiency of the spectral solver is sufficient to solve the initial data problem on a single processor, which is what we currently do.

3.4 Code Structure

We have implemented our methods in two independent codes, which we use for cross-checking of results and performance, in order to achieve reliable results from accurate and efficient simulations.

The LEAN code is based on the Cactus infrastructure [8, 55], which is a community code and thus not fully under our control. The BAM [22, 24]

code has been developed by our group and is designed to solve partial differential equations on structured meshes, in particular a coupled system of (typically hyperbolic) evolution equations and elliptic equations. The complexity of the equations is addressed by using a Mathematica package integrated into the code, which produces C-code from high-level problem descriptions in Mathematica notebooks.

The BAM code is organized as a "framework", similar in spirit to the Cactus code [8, 55], but dropping much of its complexity. The computational domain is decomposed into cubes, following standard domain-decomposition algorithms, and is parallelized with MPI. Our mesh refinement algorithm is based on the standard Berger-Oliger algorithm, but with additional buffer zones, along the lines of [42, 51].

3.5 Computational Cost

Since adaptive meshes are being used, the memory requirement is not easily expressed in terms of Cartesian box sizes. A typical large run has a cubical box as its coarsest level with about 150^3 grid points. Typical refinements use 10 levels with 1 or 2 boxes per level (depending on whether the black holes fit into one box or require two individual boxes). Depending on the refinement criterion and in particular when two boxes are merged into one larger box, the memory requirement is not constant and may exceed 150^3 by up to a factor of 4. About 100–160 grid functions are required in double precision. Extrapolating the memory requirement from the smaller runs that we performed so far, we estimate that not more than 200 GByte of RAM are needed. For 4 GByte of RAM per core on the LRZ Altix system, memory usage is therefore not a limiting factor.

However, production runs typically require on the order of 10^5 time steps on the finest level, and hence the number of floating point operations required per grid point per time step is a more stringent issue than memory. The aim is to perform simulations of black hole binaries at large separation, and the larger the separation, the longer it takes the black holes to complete one orbit. Typical large runs take from about 3 days to up to a week using checkpointing, not counting the time spent in re-queuing simulations.

Each such run results in one data point in the parameter space of black hole binaries. The number of data points in the space of initial data needed to provide useful information for gravitational wave data analysis is very much an open question right now, which we hope to address with our research.

4 Status Report of Optimization Work

4.1 General Remarks

Our BAM and LEAN codes have so far been run on Linux clusters based on AMD Opteron, and Intel architectures, and we find rather consistent perfor-

mance results for the systems we use: our in-house clusters doppler (24 AMD Opteron nodes with Infiniband network) and kepler (35 Intel Woodcrest nodes with Infiniband network), the HLRS cluster strider and the Itanium-based Altix cluster at LRZ. In summer 2005 the Jena group applied for and was granted support by the LRZ in its porting and optimization initiative for the new Altix system. We find that optimizations performed on the Altix system carry over to our in-house clusters and strider at HLRS.

Apart from standard timing tests we have used various tools: VampirTrace has been used successfully to identify communication overheads, in particular in analysis routines that compute quantities like the gravitational wave content in our simulations. For large numbers of processors, these routines have become the performance bottleneck. These calculations require computing integrals over spheres, where the collocation points are not aligned with the grid used for evolving the configuration, consequently the required interpolation operations are rather expensive in communication. Recent optimizations have removed this bottleneck, but systematic scaling results for larger number of processors are not yet available due to end of our allocation at HLRS and the downtime of HLRB2 at LRZ caused by a system-upgrade.

Various of our analysis tools require interpolation of tensor fields from Cartesian grids onto parameterized surfaces, say for wave extraction and apparent horizon finding. It was found that the scaling of the wave extraction routine in BAM was poor due to the interpolator, and this part of the code took tens of percent for typical production runs on 32 and more processors. We optimized the MPI implementation of BAM's point-wise interpolator so that wave extraction dropped to less than one percent of the overall run time.

Significant gains in efficiency have been obtained by replacing fourth order finite differencing by sixth order finite differencing as summarized in Sec. 3. This result is nontrivial for our application for several reasons: One of the problems one faces is that higher order methods generally require a larger number of buffer zones for our adaptive mesh refinement algorithm, and some experiments are required to to test different choices in practice. Furthermore, the moving puncture method deals with non-smooth features inside of the black hole horizons, and the AEI and Brownsville groups have previously reported failures to obtain stable and accurate moving puncture evolutions with sixth order methods [30, 44].

An important issue that we could identify, and which is an equally serious problem in alternative numerical relativity codes against which we benchmarked, is the overflow of registers due to the complexity of the Einstein equations. We have started to analyze the memory access pattern of the equations, and were able to perform some promising experiments, but significant further work will be required to rewrite the equations in an optimal way.

4.2 Memory Optimizations

After development of the new moving puncture methods and wave extraction tools had been completed in BAM, we optimized the memory usage of the current production code. After analyzing the use of temporary memory, it was possible to reduce the peak memory footprint counted in 3d AMR grid variables from about 280 variables to 175 variables. Part of this improvement was due to a more memory efficient implementation of the Runge-Kutta time evolution scheme.

4.3 Scaling

Our most extensive scaling tests have been performed with the second order accurate version of our code, still with non-moving boxes of refinement. In such cases we typically reach 60–90% of scaling on up to 128 processors, which has been tested on different machines. With the new fourth-order code we have so far only tested scaling on up to 32 processors. We get 77% scaling on the Cray Opteron cluster Strider going from 4 to 32 processors, and increasing the problem size by a factor close to 8. The 32-processor run corresponds to the initial $10M$ of runtime of a typical production run with 14.4 GByte memory usage (i.e. approximately 450 MByte per processor).

Already our initial scaling results on the LRZ Altix system have been very promising. In Fig. 3 we show results for the performance of the original BAM code (second order code with fixed mesh refinement). As reference point we compare to a SUN cluster (LION-XO at Penn State University) with 80 dual Opteron SunFire v20z nodes, 2.4 GHz processors, 4 GB RAM per processor, with Infiniband network, and using the gcc compiler. The timings on the first Altix test system were obtained prior to any optimizations. The test system had 64 Itanium processors at 1.6 GHz and 4 GB RAM per CPU.

The code reached 4 to 8% of peak performance or 260 to 530 MFlops on the Altix. These numbers have been improved upon in subsequent optimizations, but since the new numbers are not final yet, and also since there has been a major upgrade to the BAM code very recently, we do not want to quote premature benchmarks. However, note that even these early results show very promising network performance, see the lower panel in Fig. 3. The run labelled 5×80 requires 3.7 GB memory, run 5×128 requires 15 GB memory. For these problem sizes the Altix system clearly scales better than the Opteron/Infiniband cluster. This is important since the new moving-box/AMR algorithm of BAM and LEAN requires even more communication than the previous version of the code reported on here. We do not have recent scaling results for strider, since our grant has expired – we are currently in the process of preparing a grant application for the NEC vector machines at HLRS.

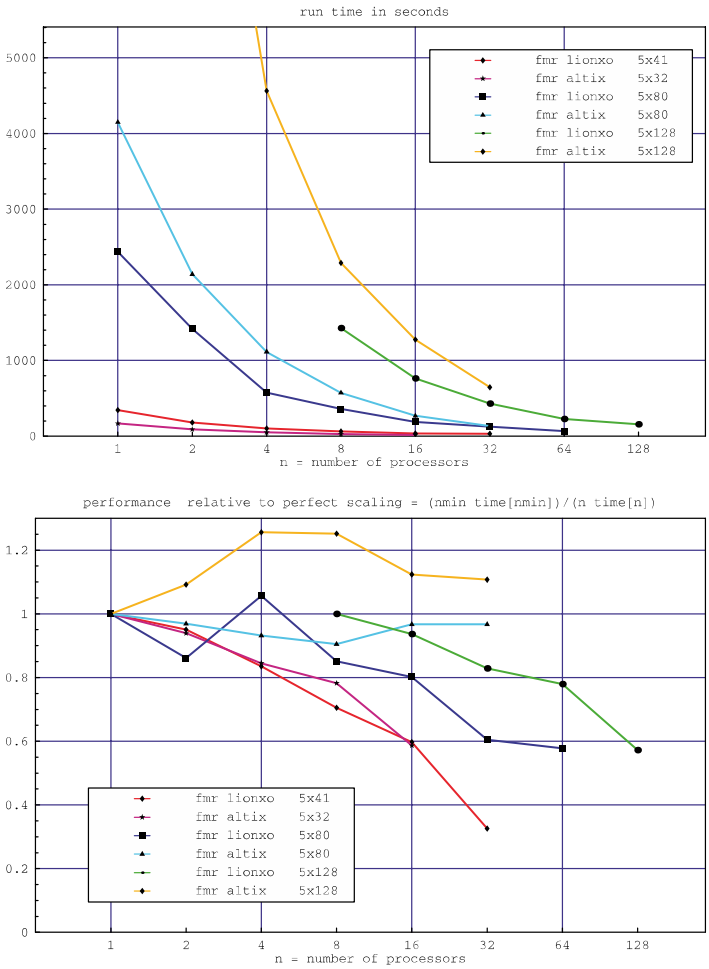

Fig. 3. Benchmarking the BAM code on the initial LRZ Altix system and an Opteron/Infiniband cluster, LION-XO. The *upper panel* shows the run times in seconds for different problem sizes, while the *lower panel* shows performance relative to perfect scaling

5 Conclusions and Plans for Future Work

We have presented a summary of recent work of the Jena numerical relativity group on simulating the coalescence of black hole binaries in full general relativity. In the last year we have obtained astrophysically relevant results: In the largest numerical parameter study to date of black-hole binaries, we have calculated the recoil of nonspinning binaries for mass ratios from 1:1 to 1:4 [35], and performed a detailed analysis of the gravitational waves emitted during unequal-mass binary mergers [17]. We have also found that extremely

large recoils are possible for spinning black holes, and estimated the kick from one configuration as 2500 km/s, large enough to eject the remnant black hole even from a giant elliptical galaxy [34]. We have also made significant progress with developing our methods and further optimizing our code. Large parameter studies forseen for the near future include the case of large mass ratios on the order of 1:10, and general spin configurations. While our initial HLRS allocation for strider has recently expired, we are currently in the process of porting our code the the NEC vector processor architecture, and prepare a grant application for the NEC SX-8.

References

1. A. A. Abramovici, W. Althouse, R. P. Drever, Y. Gursel, S. Kawamura, F. Raab, D. Shoemaker, L. Sievers, R. Spero, K. S. Thorne, R. Vogt, R. Weiss, S. Whitcomb, and M. Zuker. Ligo: The laser interferometer gravitational-wave observatory. *Science*, 256:325–333, 1992.
2. P. Ajith and *et. al.* in preparation.
3. P. Ajith and *et. al.* Phenomenological template family for black hole coalescence waveforms. In *Proceedings of the GWDAW11 conference, Potsdam, 18–21 December 2006.* in preparation.
4. Miguel Alcubierre, Gabrielle Allen, Bernd Brügmann, Edward Seidel, and Wai-Mo Suen. Towards an understanding of the stability properties of the 3+1 evolution equations in general relativity. *Phys. Rev. D*, 62:124011, 2000.
5. Miguel Alcubierre, W. Benger, B. Brügmann, G. Lanfermann, L. Nerger, E. Seidel, and R. Takahashi. 3D Grazing Collision of Two Black Holes. *Phys. Rev. Lett.*, 87:271103, 2001.
6. Miguel Alcubierre and Bernd Brügmann. Simple excision of a black hole in 3+1 numerical relativity. *Phys. Rev. D*, 63:104006, 2001.
7. Miguel Alcubierre, Bernd Brügmann, Peter Diener, Michael Koppitz, Denis Pollney, Edward Seidel, and Ryoji Takahashi. Gauge conditions for long-term numerical black hole evolutions without excision. *Phys. Rev. D*, 67:084023, 2003.
8. G. Allen, T. Goodale, J. Massó, and E. Seidel. The cactus computational toolkit and using distributed computing to collide neutron stars. In *Proceedings of Eighth IEEE International Symposium on High Performance Distributed Computing, HPDC-8, Redondo Beach, 1999.* IEEE Press, 1999.
9. Marcus Ansorg, Bernd Brügmann, and Wolfgang Tichy. A single-domain spectral method for black hole puncture data. *Phys. Rev. D*, 70:064011, 2004.
10. John Baker, Bernd Brügmann, Manuela Campanelli, Carlos O. Lousto, and Ryoji Takahashi. Plunge waveforms from inspiralling binary black holes. *Phys. Rev. Lett.*, 87:121103, 2001.
11. John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, and James van Meter. Gravitational wave extraction from an inspiraling configuration of merging black holes. 2005.
12. John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, and James van Meter. Binary black hole merger dynamics and waveforms. 2006. Unpublished, gr-qc/0602026.

13. John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, James van Meter, and M. Coleman Miller. Getting a kick out of numerical relativity. 2006. Unpublished, astro-ph/0603204.
14. Thomas W. Baumgarte and Stuart L. Shapiro. On the numerical integration of Einstein's field equations. *Phys. Rev. D*, 59:024007, 1999.
15. J. D. Bekenstein. Gravitational-Radiation Recoil and Runaway Black Holes. *Astrophys. J.*, 183:657–664, 1973.
16. Marsha J. Berger and Joseph Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.*, 53:484–512, 1984.
17. Emanuele Berti et al. Inspiral, merger and ringdown of unequal mass black hole binaries: A multipolar analysis. 2007.
18. W. B. Bonnor and M. A. Rotenberg. Transport of momentum by gravitational waves: the linear approximation. *Proc. R. Soc. Lond. A.*, 265:109–116, 1961.
19. Jeffrey M. Bowen and James W. York. Time-asymmetric initial data for black holes and black hole collisions. *Phys. Rev. D*, 21(8):2047–2056, 1980.
20. S. Brandt and B. Brügmann. A simple construction of initial data for multiple black holes. *Phys. Rev. Lett.*, 78(19):3606–3609, 1997.
21. Bernd Bruegmann et al. Calibration of moving puncture simulations. 2006.
22. B. Brügmann. Binary black hole mergers in 3D numerical relativity. *Int. J. Mod. Phys. D*, 8:85, 1999.
23. B. Brügmann, G. Gonzalez, M. Hannam, S. Husa, P. Marronetti, U. Sperhake, and W. Tichy. Gravitational wave signals from simulations of black hole dynamics. In W.E. Nagel, W. Jäger, and M. Resch, editors, *High Performance Computing in Science and Engineering '06*, pages 3–17. Springer, 2006.
24. Bernd Brügmann, Wolfgang Tichy, and Nina Jansen. Numerical simulation of orbiting black holes. *Phys. Rev. Lett.*, 92:211101, 2004.
25. Gioel Calabrese, Ian Hinder, and Sascha Husa. Numerical stability for finite difference approximations of Einstein's equations. 2005.
26. M. Campanelli, C. O. Lousto, and Y. Zlochower. Gravitational radiation from spinning-black-hole binaries: The orbital hang up. 2006.
27. Manuela Campanelli, C. O. Lousto, and Y. Zlochower. The last orbit of binary black holes. *Phys. Rev. D*, 73:061501(R), 2006.
28. Manuela Campanelli, Carlos O. Lousto, Pedro Marronetti, and Yosef Zlochower. Accurate evolutions of orbiting black-hole binaries without excision. *Phys. Rev. Letter*, 96:111101, 2006.
29. K. Danzmann. The geo project: a long baseline laser interferometer for the detection of gravitational waves. *Lecture Notes in Physics*, 410:184–209, 1992.
30. Peter Diener. private communication, 2007.
31. Éanna É. Flanagan and Scott A. Hughes. Measuring gravitational waves from binary black hole coalescence: I. signal to noise for inspiral, merger, and ringdown. *Phys. Rev. D*, 57:4535, 1998.
32. K. Gebhardt, R. Bender, G. Bower, A. Dressler, S. M. Faber, A. V. Filippenko, R. Green, C. Grillmair, L. C. Ho, J. Kormendy, T. R. Lauer, J. Magorrian, J. Pinkney, and S. Richstone, D. Tremaine. A Relationship between Nuclear Black Hole Mass and Galaxy Velocity Dispersion. *Astrophys. J*, 539:L13–L19, 2000. astro-ph/0006289.
33. GEO600 - http://www.geo600.uni-hannover.de/.
34. J. A. Gonzalez, M. D. Hannam, U. Sperhake, B. Brugmann, and S. Husa. Supermassive kicks for spinning black holes. 2007.

35. Jose A. Gonzalez, Ulrich Sperhake, Bernd Bruegmann, Mark Hannam, and Sascha Husa. Total recoil: the maximum kick from nonspinning black-hole binary inspiral. *Phys. Rev. Lett.*, 98:091101, 2007.
36. Carsten Gundlach and Jose M. Martin-Garcia. Hyperbolicity of second-order in space systems of evolution equations. 2005.
37. Mark Hannam, Sascha Husa, Bernd Bruegmann, Jose A. Gonzalez, and Ulrich Sperhake. Beyond the bowen-york extrinsic curvature for spinning black holes. 2006.
38. Mark Hannam, Sascha Husa, Niall Ó Murchadha, Bernd Brügmann, José A. González, and Ulrich Sperhake. Where do moving punctures go? *Journal of Physics: Conference series*, 2007. in press.
39. Mark Hannam, Sascha Husa, Denis Pollney, Bernd Brugmann, and Niall Ó Murchadha. Geometry and regularity of moving punctures. 2006.
40. Frank Herrmann, Deirdre Shoemaker, and Pablo Laguna. Unequal-mass binary black hole inspirals. 2006.
41. R.N. Hulse and J.H. Taylor. *Astrophys. J.*, 195:L51–L53, 1975.
42. Luis Lehner, Steven L. Liebling, and Oscar Reula. Amr, stability and higher accuracy. *Class. Quant. Grav.*, 23:S421–S446, 2006.
43. LIGO - http://www.ligo.caltech.edu/.
44. Carlos Lousto. private communication, 2006.
45. J. Magorrian, S. Tremaine, D. Richstone, R. Bender, G. Bower, A. Dressler, S. M. Faber, K. Gebhardt, R. Green, C. Grillmair, J. Kormendy, and T. Lauer. The Demography of Massive Dark Objects in Galaxy Centers. *Astron. J*, 115:2285–2305, 1998. astro-ph/9708072.
46. R. J. McLure and J. S. Dunlop. On the black hole-bulge mass relation in active ind inactive galaxies. *MNRAS*, 331:795–804, 2002. astro-ph/0108417.
47. D. Merritt and L. Ferrarese. Black hole demographics from the $m_\bullet$-σ relation. *MNRAS*, 320:L30–L34, 2001. astro-ph/0009076.
48. A. Peres. Classical Radiation Recoil. *Phys. Rev.*, 128:2471–2475, 1962.
49. Frans Pretorius. Evolution of binary black hole spacetimes. *Phys. Rev. Lett.*, 95:121101, 2005.
50. Mark A. Scheel, Harald P. Pfeiffer, Lee Lindblom, Lawrence E. Kidder, Oliver Rinne, and Saul A. Teukolsky. Solving Einstein's equations with dual coordinate frames. *Phys. Rev. D*, 74:104006, 2006.
51. Erik Schnetter, Scott H. Hawley, and Ian Hawke. Evolutions in 3D numerical relativity using fixed mesh refinement. *Class. Quantum Grav.*, 21(6):1465–1488, 21 March 2004.
52. Masaru Shibata and Takashi Nakamura. Evolution of three-dimensional gravitational waves: Harmonic slicing case. *Phys. Rev. D*, 52:5428, 1995.
53. Ulrich Sperhake. Binary black-hole evolutions of excision and puncture data. 2006.
54. J. H. Taylor and J. M. Weisberg. A new test of general relativity: Gravitational radiation and the binary pulsar PSR 1913+16. *Astrophys. J.*, 253:908–920, 1982.
55. Cactus Computational Toolkit. `http://www.cactuscode.org`.
56. VIRGO - http://www.virgo.infn.it/.
57. Y. Zlochower, J. G. Baker, M. Campanelli, and C. O. Lousto. Accurate black hole evolutions by fourth-order numerical relativity. *Phys. Rev. D*, 72:024021, 2005.

Massless Four-Loop Integrals and the Total Cross Section in $e^+ e^-$ Annihilation

J. H. Kühn, M. Steinhauser, and M. Tentyukov

Institut für Theoretische Teilchenphysik, Universität Karlsruhe, 76128 Karlsruhe, Germany

This is the report for the project `ParFORM` for the period June 2006 to June 2007.

1 Aim of the Project

The main purpose of particle physics is the explanation of the fundamental mechanism for the interaction of the elementary particles. On the experimental side the investigations take mainly place at the big accelerators at CERN (Geneva) or FERMILAB (Chicago). On the other hand it it essential to develop theoretical models which describe the fundamental interactions and which, of course, have to be confronted with the experiment.

Currently there is a well-established theory, the so-called Standard Model, which has been verified by experimental studies to a very high precision. However, there are strong hints that the Standard Model is not the final theory describing the fundamental interaction of the elementary particles. E.g., it can not describe the large amount dark matter present in the universe and it can not explain the observed huge excess of matter over anti-matter.

There are basically two routes which explore deviations from the Standard Model and which provide hints for so-called "Grand Unified Theories" (GUTs): one either performs experiments where the colliding particles have very high energy or one extracts physical observables to very high precision at lower energies and confronts them with precise theoretical calculations. Both ways are currently followed. The main emphasis of this project is centered on the second option. In particular, the corresponding calculations are important in the precise determination of the fundamental parameters of the Standard Model like coupling constants or quark masses.

The basic tool in order to perform precise calculations is perturbation theory which is based on the expansion in a small parameter, in our case the coupling constants between the various particles. In particle physics perturbation theory is organized in a very intuitive way: all mathemematical expressions

which have to be computed can be visualized in terms of so-called Feynman diagrams where each particle is represented by a line and interactions between particles by vertices. The main difficulty in practical applications is the occurrence of closed loops at higher order in perturbation theory. This project deals at the forefront of what is currently possible and considers Feynman diagrams up to five loops. (Examples will be given below.)

The basic object which is considered in this project is the photon two-point function, $\Pi(q^2)$, which is related to the correlator of two vector currents, j^μ, through[1]

$$\left(-q^2 g_{\mu\nu} + q_\mu q_\nu\right) \Pi(q^2) = i \int \mathrm{d}x \, e^{iqx} \langle 0|T j_\mu(x) j_\nu^\dagger(0)|0\rangle \,. \tag{1}$$

Some sample Feynman diagrams up to five-loop order are shown in Fig. 1. The external current j^μ is represented by a blob, the quark lines by the straight lines with an arrow and the gluons by the curly lines. Currently a closed analytic calculation taking into account the masses of the quarks, m_q, is only possible up to two loops. At three-loop order a semi-numerical method has been developed based on expansions in m_q^2/q^2 and q^2/m_q^2 which leads to accurate approximations of the three-loop result for $\Pi(q^2)$. The four-loop result for the massive and the five-loop one for the massless correlator are under investigation in this project.

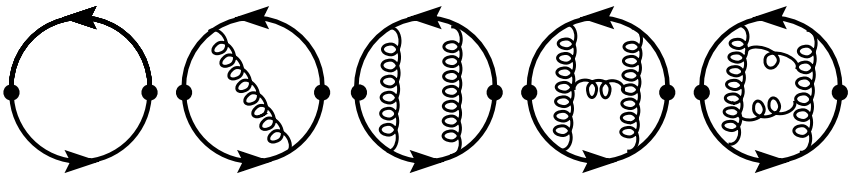

Fig. 1. Sample diagrams contributing to $\Pi(q^2)$

The imaginary part of the quantity $\Pi(q^2)$ is directly related to the physical cross section, σ, for the production of hadrons in e^+e^- annihilation, respectively, its normalized version leads to the famous R-ratio given by

$$R(s) = \frac{\sigma(e^+e^- \to \text{hadrons})}{\sigma(e^+e^- \to \mu^+\mu^-)} \;=\; 12\pi \,\mathrm{Im}\left[\Pi(q^2 = s + i\epsilon)\right] \,, \tag{2}$$

where $\sqrt{s}$ is the considered center-of-mass energy. $R(s)$ is one of the most important and clean places for precise tests of Quantum Chromodynamics (QCD) and the Standard Model. This quantity together with the related semileptonic τ lepton decay rate provide us with invaluable information about

[1] For the notation commonly used in high-energy physics we refer to the standard textbooks like, e.g., [1].

the numerical value of the strong coupling constant α_s as well as its running from the τ lepton mass to the one of the Z boson. There is also a significant amount of purely theoretical interest to higher order contributions to this quantity related, e.g., to renormalons.

Diagrammatically the imaginary part is represented by those Feynman diagrams which are obtained by cutting (i.e. separating the left and the right blob by chopping individual lines) the diagrams of Fig. 1 in all possible ways.

The main purpose is the computation of quantum corrections induced by the strong interaction to $R(s)$ which can be cast in the form

$$R(s) = \sum_{i \geq 0} \left(\frac{\alpha_s}{\pi}\right)^i \delta R^{(i)} + \text{quark mass effects}\,. \tag{3}$$

For massless quarks the corrections up to $i = 2$ are known since more than 30 years and $\delta R^{(3)}$ has been computed beginning of the nineties (for comprehensive reviews, see [2, 3]). Currently, the term $\delta R^{(4)}$ is only partially known and considered in the current project. It originates from the imaginary part of five-loop diagrams for $\Pi(q^2)$. This fact is used for the practical calculations. Actually, a sophisticated method has been developed which allows to reduce the computation of the imaginary part of five-loop diagrams to real parts of four-loop diagrams. Thus, in order to get $\delta R^{(4)}$ for massless quarks one has to consider four-loop integrals, which are actually of the same type as the ones needed for the calculation of the four-loop approximation of $\Pi(q^2)$ (in the massless case).

There are a variety of important physical applications of $\Pi(q^2)$ and its imaginary part $R(s)$. E.g., the quantity $R(s)$ can be used for the determination of α_s by confronting (3) with experimental results. In this context it is particularly promising to consider $\sqrt{s} = M_Z$, where M_Z is the mass of the Z boson, since in the past very precise data has been collected at the CERN Large Electron Positron (LEP) Collider. Furthermore, it is planned that an International Linear Collider could significantly improve the experimental data. Thus, α_s could be extracted with an uncertainty below 1% – once $\delta R^{(4)}$ is available.

The technique developed for the quantity $R(s)$ can immediately be applied to the decay rate of the τ lepton, which can be used for the determination of the strange quark mass [4]. Furthermore, if one replaces in the above discussion the vector by a scalar current, there are applications in connection to the Higgs boson. More details will be presented below.

From the above discussion it becomes clear that the main aim of this project is the computation of massless propagator-type diagrams to four-loop order within perturbative quantum field theory. The most important software tools necessary for our calculations are described in the next Section and the physical applications are discussed in Sects. 3 and 4.

2 Parallel Computer Algebra

The workhorse for most of the computations performed in this project is the computer algebra program `ParFORM` [5]. `ParFORM` is the parallel version of the Computer Algebra System `FORM` [6] which – in contrast to `Mathematica` or `Maple` – is designed for the manipulation of huge expressions ranging up to several tera byte. The latter is very crucial for for multi-loop calculations. Thus it is essential to have a fast access to the hard disk where the temporary expressions are stored.

There are essentially two implementations of `ParFORM`: one is based on MPI ("Message Passing Interface") which is quite good for systems that have processors with separated memory, and the other one is specially adopted to Symmetric-Multi-Processor (SMP) architectures with shared memory implemented using the NUMA ("Non-Uniform-Memory-Access") technology. The comparison of the speed-up curves can be seen in Fig. 2.

In the recent years `ParFORM` has been used on many different platforms ranging from SMP machines with up to 32 processors to simple PC-clusters connected by a Gigabit ethernet. In Fig. 3 the results are summarized and compared with the speed-up obtained at the XC6000 cluster. It should be stressed that, although the MPI version of `ParFORM` has to be used on the XC6000 cluster the same speed-up as for the faster SMP version (cf. Fig. 2) has been obtained.

The run-time of our problems varies from a few days or weeks up to about two months. Due to the very structure of `FORM` and due to the organization of the calculation it is not possible to set check points. Since the maximum CPU time at the Landeshöchstleistungsrechner is limited to about seven days it is only possible to submit small and medium-sized jobs.

In Fig. 4 the performance of `ParFORM` is shown for a typical job where up to 60 processors have been used. For the XC6000 a good scaling behaviour is observed up to about 16 processors. Above approximately 24 processors the saturation region starts and only a marginal gain is observed once 60 processors are employed. The situation is much worse for the new Lan-

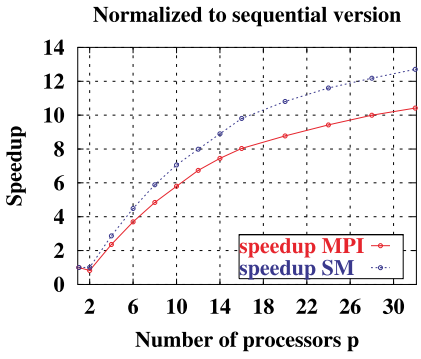

Fig. 2. Speed-up as a function of the used processors. The *full* (*red*) *curve* is based on the MPI and the *dotted* (*blue*) *curve* on the shared-memory (SM) version of `ParFORM`. The data point have been obtained on a 32-processor Itanium computer

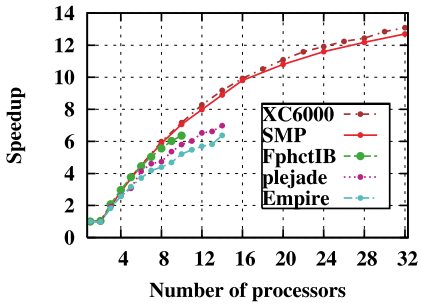

Fig. 3. Speed-up curves for various computer clusters with different interconnections

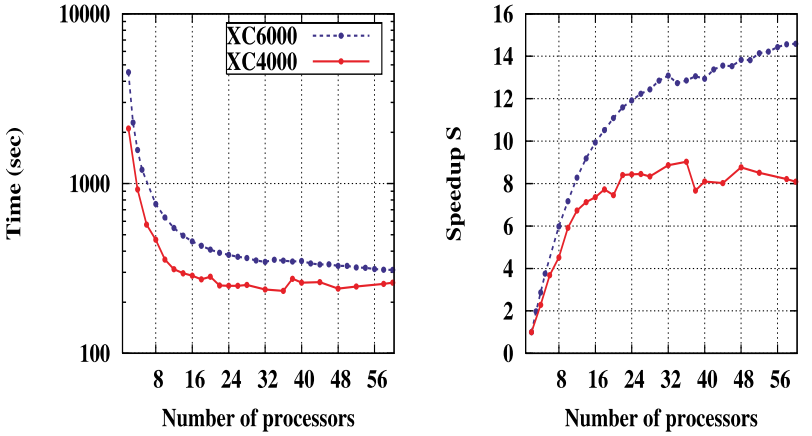

Fig. 4. CPU-time and speedup curve for a typical job on the XC6000 and XC4000

deshöchstleistungsrechner (XC4000). Beyond about ten processors the system is very unstable and is thus less attractive for our applications. It seems that the interconnection of the individual nodes is much worse than for the XC6000 cluster – at least for our applications.

Very recently a different concept for the parallelization of `FORM` has been developed. The basic idea is the use of POSIX[2] threads in order to realize the communication between the various processors of a shared memory machine. The main application is thus centered around multi-core machines with two, four or eight cores. First tests of `TFORM` [7] were quite successful and comparable speed-up to `ParFORM` could be achieved. In future we will continue to further development of `TFORM`. Due to the hardware structure of the Landeshöchstleistungsrechner this is very promising. It is particularly tempting to combine `TFORM` and `ParFORM` in order to reach an optimal speed-up.

A further program, `Crusher`, which has been developed in our group, implements the so-called Laporta-algorithm [8, 9]. In high-energy physics the Laporta-algorithm is a widespread tool which is used to reduce the huge num-

[2] "a Portable Operating System Interface for uniX"

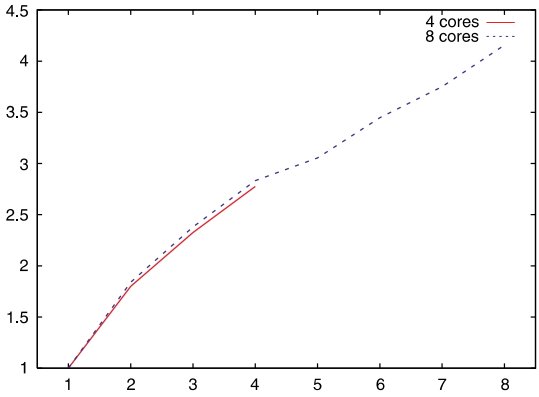

Fig. 5. Speed-up curve for the program `TCrusher`. The *solid* (*red*) and *dotted* (*blue*) *curves* corresponds to computers with two and four dual core opteron CPUs, respectively

ber of Feynman integrals occuring in the calculation of physical quantity to a small set of basis integrals. The reduction is essentially based on a particular implementation of the Gauss elimination algorithm which is applied to a system of equation obtained from the original Feynman diagrams. The main problem in the practical implementation of the Laporta-algorithm is that the number of equations ranges typically up to several millions and thus an effective method is mandatory. `Crusher` is written in `C++` and uses the computer algebra program `Fermat` [10] for the manipulation of the coefficients of the sought integrals which are rational functions in the space-time, mass and momentum of the space-time and mass parameters of the problem. Very recently a thread-based version of `Crusher` has been developed, `TCrusher`. In Fig. 5 the speed-up curves obtained on a four- and eight-core opteron computer, are shown as a function of the number of used threads. The speed-up curves have been generated for a relatively small problem with a runtime of approximately one hour. They show the speed-up for the main and most time consuming part of the program, namely the solving of the equations. A speed-up of three can be reached with four threads. For problems with a higher complexity an even larger speed-up can be reached.

3 Massless Four-Loop Integrals: $\sigma(e^+e^- \to \text{hadrons})$

For most of the compute-jobs connected to the massless four-loop integrals we have used about eight processors which leads to the total amount of about 50 processors assuming six jobs in the batch queue. For some jobs it turned out that it is advantageous to use more processors up to a maximum of 32. Almost all results connected to this sub-project were obtained on the XC6000 Itanium nodes.

As mentioned in the Introduction $R(s)$ is currently known to order α_s^3 and the corresponding theoretical uncertainty in the value of $\alpha_s(M_Z)$ is around 3% which is the same as the experimental one. In future the experimental errors

will be decreased and thus it is necessary to compute the $\mathcal{O}(\alpha_s^4)$ term in order to reduce also the uncertainties in the theoretical prediction. The calculation is highly non-trivial and requires a lot of preparation work to be done which is necessary to fine tune the programs and to accumulate experience.

The preparation work has been under way during last years. A few related projects have been successfully completed and their results have been published [11, 12, 13, 14]. Thus, the theoretical possibility of the complete calculation has been demonstrated.

As mentioned in Sect. 1 the order α_s^4 contribution to $R(s)$ is related to the absorptive part of the five-loop vector current correlator, whose calculation eventually boils down to finding a host of four-loop propagator-type integrals (p-integrals).

In order to cope with the problem the special package, `BAICER`, has been created. This is a `FORM` package capable to analytically compute p-integrals up to four loops. Since the conventional method [15] to treat such integrals is not applicable (and is commonly considered as not being possible for the foreseeable future), `BAICER` uses a completely different approach [16, 17, 18]. Namely, the package computes coefficients in decomposition of a given p-integral into the irreducible ones. The coefficients are known to be rational functions of the space-time dimension D and are computed as expansion over $1/D$ as $D \to \infty$. From the knowledge of sufficiently many terms in the expansion one can reconstruct their exact form. The terms in the $1/D$ expansion are expressed in terms of simple Gaussian integrals. Starting from four loops very large number of the latter are necessary to calculate. For example, the first non-trivial five-loop QCD result obtained with BAICER – the $\alpha_s^3 n_f^2$ contribution to $R(s)$ in QCD [19] – involves several billions of Gaussian integrals.

During 2006 and 2007 the following problems have been computed with `BAICER` on the basis of our local SGI multi-processor computer and the XC6000 in the Rechenzentrum.

1. The $\mathcal{O}(\alpha_s^5)$ correction to the $H \to gg$ partial width of the Standard Model Higgs boson with intermediate mass $M_H < 2M_t$ [20]. Its knowledge is useful because with around 20% the $\mathcal{O}(\alpha_s^4)$ correction is sizeable. The new four-loop correction increases the total Higgs-boson hadronic width by a small amount of order 1 per mille and stabilizes significantly the residual scale dependence.
2. The four-loop anomalous dimensions of the tensor current and the non-singlet twist-two operator. In the second case we do not only compute the anomalous dimension but also the four-loop matrix element of the (zero-momentum) insertion of the operator between two quark states. Such information is useful for lattice calculations within so-called regularization-independent renormalization schemes [21, 22].
   The knowledge of the anomalous dimension of the tensor quark current is important because it occurs in the matching between currents in QCD effective quark theory [23]. It also appears during the extraction of one of

the least known CKM parameters, $|V_{ub}|$, from the corresponding leptonic and semileptoinc decays of B-mesons (see, e.g., [24, 25]).
The results for specific (and phenomenologically most interesting) value for the number of the active quark flavours, $n_f = 3$, have been published in [26]. Meanwhile we have finished calculations for generic n_f; the corresponding publication in preparation.

3. The most extensive calculation performed on the XC6000 is directly related to R(s); namely the evaluation of the β-function for the so-called Quenched[3] QED (QQED) in the five loop approximation. In fact, the quantity can be considered as an important gauge independent contribution to R(s) at order α_s^4 (namely the one proportional to the colour structure C_F^4). The QQED β-function is also a fascinating theoretical object by its own right:
   (i) It is scheme independent in all orders.
   (ii) Its coefficients are simple rational numbers at one, two, three and four loops: $(4/3, 4, -2, -46)$.
   (iii) There is a belief that this characteristic reflects some deep not yet understood property of the quantum field theory which should be valid in all orders (see, e.g., [27, 28]).

   The result of our calculation for the five-loop term of the QQED β-function unexpectedly contains an irrational contribution proportional to $\zeta(3)$. At the moment we are checking our calculations and simultaneously we extend them to include all other colour structures necessary for $R(s)$.

4 Massive Vacuum Integrals: $\Pi(q^2)$ to Four Loops

As mentioned in the Introduction, it is also one of the aims of the project to compute four-loop corrections to the polarization function $\Pi(q^2)$ including massive quarks. An immediate application is discussed in [29] where very precise values for the charm and bottom quark mass have been extracted. Furthermore, after taking the imaginary part, one would obtain the full quark mass dependence of $R(s)$ to order α_s^3.

The main challenge in the computation of $\Pi(q^2)$ is due the occurrence of two mass scales: the quark mass and the external momentum. After writing down the mathematical expression for the four-loop Feynman diagrams contributing to $\Pi(q^2)$ one has millions of different integrals. As mentioned in Sect. 1, the standard approach is to reduce them to a small set of basis integrals with the help of the Laporta-algorithm which we implemented in the parallel program `TCrusher`.

The status of this sub-project is as follows: `TCrusher` has successfully been applied to the three-loop diagrams. The reduction to master integrals took

[3] i.e., diagrams with closed fermion loops are not considered

about 48 CPU hours where 4 processors have been used. Once the result is expressed in terms of basis integrals the latter have to be computed. A promising approach is based on an expansion for small external momentum. For illustration we present the results for the coefficient of the 30th expansion term of $\Pi(q^2)$ in limit of small external momentum:

$$\Pi(q^2)|_{3\,\mathrm{loops}} = \ldots + \frac{\alpha_s}{\pi}\left(\frac{q^2}{4m_Q^2}\right)^{30}\Bigg\{ - 0.88385 C_A\, C_F + 6.0034 C_F^2 - 0.092427 C_F T_F + 0.52365 C_F n_l T_F \Bigg\} + \ldots\,, \tag{4}$$

($n_l = n_f - 1$) which is in agreement with the literature [30]. At this point is should be stressed that in the approach where in a first step an expansion for small q is done and afterwards the reduction to the basis integrals is performed it is hardly possible to reach the tenth expansion term [31].

Currently the reduction of the four-loop diagrams is performed on the cluster XC4000. On average 64 processors are used. The average CPU time needed for the reduction of a typical diagram amounts to 100 days.

Acknowledgements

The computations presented in this contribution were performed on the Landeshöchstleistungsrechner XC4000 and XC6000.

References

1. F. J. Yndurain, "Relativistic quantum mechanics and introduction to field theory," *Berlin, Germany: Springer (1996) 332 p. (Text and monographs in physics)*
2. K. G. Chetyrkin, J. H. Kuhn and A. Kwiatkowski, "QCD corrections to the e+ e- cross-section and the Z boson decay rate: Concepts and results," Phys. Rept. **277** (1996) 189.
3. R. V. Harlander and M. Steinhauser, "rhad: A program for the evaluation of the hadronic R-ratio in the perturbative regime of QCD," Comput. Phys. Commun. **153** (2003) 244 [arXiv:hep-ph/0212294].
4. K. G. Chetyrkin and A. Khodjamirian, "Strange quark mass from pseudoscalar sum rule with O(alpha(s)**4) accuracy," Eur. Phys. J. C **46** (2006) 721 [arXiv:hep-ph/0512295].
5. M. Tentyukov, D. Fliegner, M. Frank, A. Onischenko, A. Retey, H. M. Staudenmaier and J. A. M. Vermaseren, "ParFORM: Parallel Version of the Symbolic Manipulation Program FORM," arXiv:cs.sc/0407066;
   M. Tentyukov, H. M. Staudenmaier and J. A. M. Vermaseren, "ParFORM: Recent development," Nucl. Instrum. Meth. A **559** (2006) 224.
   H. M. Staudenmaier, M. Steinhauser, M. Tentyukov, J. A. M. Vermaseren, "ParFORM," Computeralgebra Rundbriefe **39** (2006) 19.
   See also http://www-ttp.physik.uni-karlsruhe.de/∼parform.

6. FORM version 3.0 is described in: J. A. M. Vermaseren, "New features of FORM," arXiv:math-ph/0010025;
for recent developments, see also: M. Tentyukov and J. A. M. Vermaseren, "Extension of the functionality of the symbolic program FORM by external software," arXiv:cs.sc/0604052;
FORM can be obtained from the distribution site at http://www.nikhef.nl/∼form.
7. M. Tentyukov and J. A. M. Vermaseren, "The multithreaded version of FORM," arXiv:hep-ph/0702279.
8. S. Laporta and E. Remiddi, "The analytical value of the electron (g-2) at order alpha 3 in QED," Phys. Lett. B **379** (1996) 283 [arXiv:hep-ph/9602417].
9. S. Laporta, "High-precision calculation of multi-loop Feynman integrals by difference equations," Int. J. Mod. Phys. A **15**, 5087 (2000) [arXiv:hep-ph/0102033].
10. R. H. Lewis, Fermat's User Guide, http://www.bway.net/˜lewis.
11. P. A. Baikov, K. G. Chetyrkin and J. H. Kuhn, "Five-loop vacuum polarization in pQCD: O(alpha(s)**4 N(f)**2) results," Nucl. Phys. Proc. Suppl. **116** (2003) 78.
12. P. A. Baikov, K. G. Chetyrkin and J. H. Kuhn, "Vacuum polarization in pQCD: First complete O(alpha(s)**4) result," Nucl. Phys. Proc. Suppl. **135** (2004) 243.
13. P. A. Baikov, K. G. Chetyrkin and J. H. Kuhn, "Strange quark mass from tau lepton decays with O(alpha(s**3)) accuracy," Phys. Rev. Lett. **95** (2005) 012003 [arXiv:hep-ph/0412350].
14. P. A. Baikov, K. G. Chetyrkin and J. H. Kuhn, "Scalar correlator at O(alpha(s)**4), Higgs decay into b-quarks and bounds on the light quark masses," Phys. Rev. Lett. **96** (2006) 012003 [arXiv:hep-ph/0511063].
15. K. G. Chetyrkin and F. V. Tkachov, "Integration By Parts: The Algorithm To Calculate Beta Functions In 4 Loops," Nucl. Phys. B **192** (1981) 159.
16. P. A. Baikov, "Explicit solutions of the multi-loop integral recurrence relations and its application," Nucl. Instrum. Meth. A **389** (1997) 347 [arXiv:hep-ph/9611449].
17. P.A. Baikov, "Explicit solutions of the three loop vacuum integral recurrence relations," Phys. Lett. B **385** (1996) 404 [arXiv:hep-ph/9603267].
18. P. A. Baikov, "The criterion of irreducibility of multi-loop Feynman integrals," Phys. Lett. B **474** (2000) 385 [arXiv:hep-ph/9912421].
19. P. A. Baikov, K. G. Chetyrkin and J. H. Kuhn, "The cross section of e+ e- annihilation into hadrons of order alpha(s)**4 n(f)**2 in perturbative QCD," Phys. Rev. Lett. **88** (2002) 012001 [arXiv:hep-ph/0108197].
20. P. A. Baikov and K. G. Chetyrkin, "Higgs decay into hadrons to order alpha(s)**5," Phys. Rev. Lett. **97** (2006) 061803 [arXiv:hep-ph/0604194].
21. G. Martinelli, C. Pittori, C. T. Sachrajda, M. Testa and A. Vladikas, "A General Method For Nonperturbative Renormalization Of Lattice Operators," Nucl. Phys. B **445** (1995) 81 [arXiv:hep-lat/9411010].
22. M. Gockeler *et al.*, "Nonperturbative renormalisation of composite operators in lattice QCD," Nucl. Phys. B **544** (1999) 699 [arXiv:hep-lat/9807044].
23. D. J. Broadhurst and A. G. Grozin, "Matching QCD And Hqet Heavy - Light Currents At Two Loops And Beyond," Phys. Rev. D **52** (1995) 4082 [arXiv:hep-ph/9410240].
24. V. M. Braun, T. Burch, C. Gattringer, M. Gockeler, G. Lacagnina, S. Schaefer and A. Schafer, "A lattice calculation of vector meson couplings to the vector

and tensor currents using chirally improved fermions," Phys. Rev. D **68** (2003) 054501 [arXiv:hep-lat/0306006].
25. D. Becirevic, V. Lubicz, F. Mescia and C. Tarantino, "Coupling of the light vector meson to the vector and to the tensor current," JHEP **0305** (2003) 007 [arXiv:hep-lat/0301020].
26. P. A. Baikov and K. G. Chetyrkin, "New four loop results in QCD," Nucl. Phys. Proc. Suppl. **160** (2006) 76.
27. D. J. Broadhurst, "Four-loop Dyson-Schwinger-Johnson anatomy," Phys. Lett. B **466** (1999) 319 [arXiv:hep-ph/9909336].
28. A. Connes and D. Kreimer, "Renormalization in quantum field theory and the Riemann-Hilbert problem," JHEP **9909** (1999) 024 [arXiv:hep-th/9909126].
29. J. H. Kuhn, M. Steinhauser and C. Sturm, "Heavy quark masses from sum rules in four-loop approximation," arXiv:hep-ph/0702103.
30. R. Boughezal, M. Czakon and T. Schutzmeier, "Four-loop tadpoles: Applications in QCD," Nucl. Phys. Proc. Suppl. **160** (2006) 160 [arXiv:hep-ph/0607141].
31. K. G. Chetyrkin, J. H. Kuhn and M. Steinhauser, "Heavy quark current correlators to O(alpha(s)**2)," Nucl. Phys. B **505** (1997) 40 [arXiv:hep-ph/9705254].

Structural Transitions in Colloidal Suspensions

Martin Hecht and Jens Harting

Institut für Computerphysik, Pfaffenwaldring 27, 70569 Stuttgart, Germany

Summary. In suspensions of colloidal particles different types of interactions are in a subtle interplay. In this report we are interested in sub-micro meter sized Al_2O_3 particles which are suspended in water. Their interactions can be adjusted by tuning the *p*H-value and the salt concentration. In this manner different microscopic structures can be obtained. Industrial processes for the production of ceramics can be optimized by taking advantage of specific changes of the microscopic structure. To investigate the influences of the *p*H-value and the salt concentration on the microscopic structure and the properties of the suspension, we have developed a coupled Stochastic Rotation Dynamics (SRD) and Molecular Dynamics (MD) simulation code. The code has been parallelized using MPI. We utilize the pair correlation function and the structure factor to analyze the structure of the suspension. The results are summarized in a stability diagram. For selected conditions we study the process of cluster formation in large scale simulations of dilute suspensions.

Key words: Stochastic Rotation Dynamics; Molecular Dynamics; colloids; clustering

1 Introduction

Colloid science is a very fascinating research field, gaining more and more importance in the last years. It closely connects physics, chemistry, material science, biology, and several branches of engineering technology. According to its key role in modern science a considerable amount of research has been performed to describe colloidal suspensions from a theoretical point of view and by simulations [16, 28, 29, 41, 47, 49] as well as to understand the particle-particle interactions [3, 11, 12, 15, 51, 52], the phase behavior [10, 23, 32, 50], the relevant processes on the microscale and their influence on macroscopic parameters [13, 40, 54]. Colloidal suspensions are in fact complicated systems, since depending on the particle sizes, materials, and concentrations, different interactions are of relevance and often several of them are in a subtle interplay: electrostatic repulsion, depletion forces, van der Waals attraction,

hydrodynamic interaction, Brownian motion, and gravity are the most important influences. The properties of the suspension strongly depend on the balance of the microscopic forces between the particles. Especially for industrial processes, where one needs to optimize certain material properties a detailed understanding of the relevant influences is needed. The stability of different microstructures and especially the clustering process are key properties which are of interest.

In our work we investigate these properties, focusing on Al_2O_3 particles suspended in water. This is a widely used material in ceramics [37]. We have developed a simulation code for a Brownian suspension [20] and have adjusted the simulation parameters so that the simulation corresponds quantitatively to a real suspension such that experimental data can be compared directly. The diffusion coefficient, sedimentation velocity [20], and the viscosity of the suspension can be reproduced [17]. We also have tested the influence of polydispersity and found that its influence on the results is small. It is much more important to choose the correct mean size of the particles [17]. For Al_2O_3 suspensions attractive van der Waals forces are important for the behavior of this material. Electrostatic repulsion of the charged particles counteracts the attraction and can prevent clustering depending on the particle surface charge. In [17] we have presented how one can relate parameters of DLVO potentials [11, 52] with experimental conditions. In the experiment one can control the pH-value and the salt concentration. The latter can be expressed by the ionic strength I, which is an effective concentration of all ions present in the solution. Both, the pH-value and the ionic strength, influence the charge of the colloidal particles. We have shown that for not too strongly attractive forces one can obtain reasonable quantitative agreement with experimental results.

Three regimes can be identified and plotted in a stability diagram [17], which we want to investigate here in more detail: A clustered regime, in which particles aggregate to clusters, a fluid-like and stable stable suspension and a repulsive region, for which the microstructure is similar to the ones known from glassy systems. From our previous work we know that our model works well, even quantitatively, in the suspended regime of the stability diagram and close to the borders between the different microstructures. Here we extend our investigations to different pH-values, deeper in the clustered regime, and to the repulsive structure. We expect to gain insight to the microscopic structure on a qualitative level.

On these grounds we have explored the stability diagram of Al_2O_3 suspensions. The particles are uncharged close to the so called "isoelectric point" at $p\mathrm{H} = 8.7$. There, for all ionic strengths the particles form clusters. For lower pH-values particles can be stabilized in solution by the electrostatic repulsion due to the charge the particles carry in this case. For low pH-values, low salt concentrations, and high volume fractions a repulsive structure can be found.

In the following section we shortly describe our simulation method. After that we discuss the properties which can be found in our suspensions and how

different regimes of the stability diagram are distinguished. In the section thereafter we describe our simulation setup. Then, we present our results and discuss the criteria we apply to characterize the microstructures. We utilize the pair correlation function and the structure factor to characterize the clustering behavior. Both of them in principle contain the same information, but we concentrate on certain peaks in either of them. Each peak in the correlation function and in the structure factor corresponds to a certain length scale and we chose either the correlation function or the structure factor, depending on which of the two quantities is more suitable under numerical criterions to observe on a given length scale. To characterize the repulsive region we evaluate the mean squared displacement (MSD), which shows a plateau, if the particle motion consists of different processes acting on well separated time scales. Finally, the results are summarized in a stability diagram for our Al_2O_3-suspension. It shows the behavior of the suspension in an intuitive way and helps to design industrial processes using this material. After that, we turn to dilute suspensions of only 5% volume fraction and study cluster growth at low shear rates in these suspensions. Finally, we shortly summarize our results. The results which we present in this report have been accepted for publication in [18, 19].

2 Simulation Method

Our simulation method is described in detail in [17, 20] and consists of two parts: a Molecular Dynamics (MD) code, which treats the colloidal particles, and a Stochastic Rotation Dynamics (SRD) simulation for the fluid solvent. In the MD part we include effective electrostatic interactions and van der Waals attraction, known as DLVO potentials [11, 52], a lubrication force and Hertzian contact forces. DLVO potentials are composed of two terms, the first one being an exponentially screened Coulomb potential due to the surface charge of the suspended particles

$$V_{\mathrm{Coul}} = \pi\varepsilon_r\varepsilon_0 \left[\frac{2+\kappa d}{1+\kappa d}\cdot\frac{4k_{\mathrm{B}}T}{ze}\tanh\left(\frac{ze\zeta}{4k_{\mathrm{B}}T}\right)\right]^2 \times \frac{d^2}{r}\exp(-\kappa[r-d]), \quad (1)$$

where d denotes the particle diameter, r the distance between the particle centers, e the elementary charge, T the temperature, k_{B} the Boltzmann constant, and z is the valency of the ions of added salt. ε_0 is the permittivity of the vacuum, $\varepsilon_r = 81$ the relative dielectric constant of the solvent, κ the inverse Debye length defined by $\kappa^2 = 8\pi\ell_B I$, with ionic strength I and Bjerrum length $\ell_B = 7$. The effective surface potential ζ can be related to the pH-value of the solvent with a $2pK$ charge regulation model [17]. The Coulomb term competes with the attractive van der Waals interaction ($A_{\mathrm{H}} = 4.76 \cdot 10^{-20}$ J is the Hamaker constant) [21]

$$V_{\mathrm{VdW}} = -\frac{A_{\mathrm{H}}}{12}\left[\frac{d^2}{r^2-d^2} + \frac{d^2}{r^2} \;+2\ln\left(\frac{r^2-d^2}{r^2}\right)\right]. \quad (2)$$

The attractive contribution V_{VdW} is responsible for the cluster formation we observe. However, depending on the pH-value and the ionic strength, it may be overcompensated by the electrostatic repulsion. When particles get in contact, the potential has a minimum. However, (2) diverges due to the limitations of DLVO theory. We cut off the DLVO potentials and model the minimum by a parabola as described in [20]. The particle contacts are modeled as Hertzian contacts and for non-touching particles. Below the resolution of the SRD algorithm short range hydrodynamics is corrected by a lubrication force, which we apply within the MD framework, as we have explained in [17, 20]. For the integration of translational motion of the colloidal particles we utilize a velocity Verlet algorithm [5].

For the simulation of a fluid solvent, many different simulation methods have been proposed: Stokesian Dynamics (SD) [6, 7, 41], Accelerated Stokesian Dynamics (ASD) [45, 46], pair drag simulations [47], Brownian Dynamics (BD) [21, 22], Lattice Boltzmann (LB) [27, 28, 29, 30], and Stochastic Rotation Dynamics (SRD) [20, 26, 38]. These mesoscopic fluid simulation methods have in common that they make certain approximations to reduce the computational effort. Some of them include thermal noise intrinsically, or it can be included consistently. They scale differently with the number of embedded particles and the complexity of the algorithm differs largely.

We apply the Stochastic Rotation Dynamics method (SRD) introduced by Malevanets and Kapral [33, 34]. It intrinsically contains fluctuations, is easy to implement, and has been shown to be well suitable for simulations of colloidal and polymer suspensions [4, 17, 20, 26, 38, 42, 53] and recently for star-polymers in shear flow [44]. The method is also known as "Real-coded Lattice Gas" [26] or as "multi-particle-collision dynamics" (MPCD) [43]. It is based on so-called fluid particles with continuous positions and velocities. A streaming step and an interaction step are performed alternately. In the streaming step, each particle i is moved according to

$$\mathbf{r}_i(t+\tau) = \mathbf{r}_i(t) + \tau\ \mathbf{v}_i(t), \tag{3}$$

where $\mathbf{r}_i(t)$ denotes the position of the particle i at time t and τ is the time step. In the interaction step the fluid particles are sorted into cubic cells of a regular lattice and only the particles within the same cell interact among each other according to an artificial interaction. The interaction step is designed to exchange momentum among the particles, but at the same time to conserve total energy and total momentum within each cell, and to be very simple, i.e., computationally cheap: each cell j is treated independently. First, the mean velocity $\mathbf{u}_j(t') = \frac{1}{N_j(t')} \sum_{i=1}^{N_j(t')} \mathbf{v}_i(t)$ is calculated. $N_j(t')$ is the number of fluid particles contained in cell j at time $t' = t + \tau$. Then, the velocities of each fluid particle in cell j are rotated according to

$$\mathbf{v}_i(t+\tau) = \mathbf{u}_j(t') + \mathbf{\Omega}_j(t') \cdot [\mathbf{v}_i(t) - \mathbf{u}_j(t')]. \tag{4}$$

$\mathbf{\Omega}_j(t')$ is a rotation matrix, which is independently chosen at random for each time step and each cell. We use rotations about one of the coordinate axes

by an angle $\pm\alpha$, with α fixed. The coordinate axis as well as the sign of the rotation are chosen at random, resulting in 6 possible rotation matrices. To remove anomalies introduced by the regular grid, one can either choose a mean free path of the order of the cell size or shift the whole grid by a random vector once per SRD time step as proposed by Ihle and Kroll [24, 25].

Three different methods to couple the SRD and the MD simulation have been introduced in the literature. Inoue *et al.* proposed a way to implement no slip boundary conditions on the particle surface [26]. Padding and Louis very recently came up with full slip boundaries, where the fluid particles interact via Lennard-Jones potentials with the colloidal particles [39]. Falck *et al.* [14] have developed a "more coarse grained" method which we use for our simulations and which we descibe shortly in the following.

To couple the colloidal particles to the fluid, the colloidal particles are sorted into the SRD cells and their velocities are included in the rotation step. One has to use the mass of each particle –colloidal or fluid particle –as a weight factor when calculating the mean velocity

$$\mathbf{u}_j(t') = \frac{1}{M_j(t')} \sum_{i=1}^{N_j(t')} \mathbf{v}_i(t) m_i, \tag{5}$$

$$\text{with} \qquad M_j(t') = \sum_{i=1}^{N_j(t')} m_i, \tag{6}$$

where we sum over all colloidal and fluid particles in the cell, so that $N_j(t')$ is the total number of both particles, fluid plus colloidal ones. m_k is the mass of the particle with index i and $M_j(t')$ gives the total mass contained in cell j at time $t' = t + \tau$. To some of our simulations we apply shear. This is realized by explicitly setting the mean velocity $\mathbf{u}_j$ to the shear velocity in the cells close to the border of the system. Both, colloidal and fluid particles, are involved in this additional step. A thermostat is applied to remove the energy introduced to the system by the shear force. We have described the simulation method in more detail in [17, 20].

A single simulation run as presented in these papers took between one and seven days on a 3GHz Pentium CPU. However, for strongly clustering systems we easily end up with only a single cluster inside the simulation volume. In order to be able to gather statistics on cluster growth and formation, as well as to minimize finite size effects, we parallelized our code. While MD codes have been parallelized by many groups, only few parallel implementations of a coupled MD and SRD program exist. This is in contrast to the number of parallel implementations of other mesoscopic simulation methods like for example the lattice Boltzmann method. A possible explanation is that SRD is a more recent and so far not as widely used algorithm causing the parallelization to be a more challenging task.

We utilize the Message Pasing Interface (MPI) to create a C++ code based on domain decomposition for both involved simulation methods. In

the MD code the position of neighbouring particles is needed to compute the interactions. Since the intractions have a limited range, and a linked cell algorithm is already used in the serial code, we apply linked cells here as well. Particle positions at the border of the domain of each processor are communicated to the neighbouring processors for the calculation of the forces. Then, the propagation step is performed and particle positions are updated, whereby the particles crossing a domain boundary are transferred from one processor to the other one.

Since (in principle), fluid particles can travel arbitrary large distances in one time step, one either has to limit the distance they can move, or one needs all-to-all communication between the processors. Even though the mean free path in our systems is small enough to limit communication to nearest neighbours only, the current version of our code tries to be as general as possible and allows fluid particles to move to any possible position in the total simulation volume within a single timestep. First, we determine locally which fluid particles have to be sent to which destination CPU and collect all particles to be sent to the same destination into a single MPI message. If no particles are to be sent, a zero dummy message is transmitted. On the receiving side, `MPI_Probe` with the `MPI_ANY_SOURCE` option is utilized to determine the sender's rank and the number of particles to be accomodated. Now, `MPI_Recv` can be used to actually receive the message. All processors send and receive in arbitrary order, thus waiting times are kept at a minimum allowing a very efficient communication. The standard MPI all-to-all communication procedure should be less efficient since the size of every message would be given by the size of the largest message. However, we still do find a substantial communication overhead from our benchmark tests of the scalability of the code. Due to this overhead, we are currently limited to 32 CPUs. In order to achieve Gallilean invariance, a random shift of the SRD lattice is performed for every rotation step [24, 25]. Since the domains managed by each CPU do not move, this would include the borders between the processors to cross SRD cells, which is undesirable. Therefore, we keep the position of the lattice fixed and shift the fluid particle positions before sorting them into the cells instead. After the rotation step they are shifted back.

3 Background

We examine the microstructures obtained in our simulations for different conditions. We vary the pH-value and the ionic strength I. The shear rate $\dot{\gamma}$ as an external influence is varied as well. We classify the microstructures in three categories: suspended, clustered, and repulsive. In the suspended case, the particles can move freely in the fluid and do not form stable clusters. In the clustered regime the particles form clusters due to attractive van der Waals forces. These clusters can be teared apart if shear is applied. In some of our simulations the clusters are very weakly connected and at small shear rates

they are not only broken up into smaller pieces, but they dissolve to freely moving individual particles. In this case, we assign the microstructure to the suspended region, although in complete absence of the shear flow clusters are formed. At the borders between the different regimes in fact no sharp transitions can be observed. The DLVO forces rather steadily increase and compete with the hydrodynamic interactions. Accordingly, in experiments one cannot observe a sudden solidification, but a steadily increasing viscosity when leaving the suspended regime [17].

Similarly as for attractive forces, repulsive interactions can restrict the mobility of the particles. If this happens, the mean squared displacement of the particles shows a pronounced plateau, as it can be found in glassy systems. However, we speak of a "repulsive structure", because the change of the viscosity is not as strong as in glasses, where it often changes by many oders of magnitude, when the glass transition is approached. In addition, to claim a system shows a glassy behavior would require to investigate the temperature dependence of a typical time (e.g. particle diffusion time) and to show its divergence as the glass temperature is approached. This is difficult to do in the framework of our simulation model [20] and therefore we prefer to speak about a "repulsive structure" which might be identified as a colloidal glass in future work.

Here we would like to emphasize the analysis of the microstructure for different conditions. Our aim is to reproduce a so-called stability diagram by simulations. The stability diagram depicts the respective microstructure depending on the *p*H-value and the ionic strength I. We apply different numerical tools to analyze the microstructure in our simulations and finally arrive at a stability diagram shown in Fig. 7, which summarizes the results which we present in the following sections.

4 Simulation Setup

In this study the colloidal particles are represented by three dimensional spheres of $d = 0.37\,\mu\text{m}$ in diameter. This is the mean diameter of the particles used in the experiments to which we refer in [17]. We have simulated a small volume, $24\,d = 8.88\,\mu\text{m}$ long in x-direction, which is the shear direction, and $12\,d = 4.44\,\mu\text{m}$ long in y- and z-direction. We have varied the volume fraction between $\Phi = 10\,\%$ (660 particles) and $\Phi = 40\,\%$ (2640 particles). Most of the simulations were performed at $\Phi = 35\,\%$ (2310 particles). To study low volume fractions $\Phi = 5\,\%$ we have enlarged the simulation volume to $24\,d = 8.88\,\mu\text{m}$ in each direction and we have further scaled up the system in each dimesion by a factor of 2 or 4, resulting in a cube of $48d = 17.76\mu\text{m}$. For selected *p*H-values and ionic strengths we have studied the cluster growth of dilute suspensions at low shear rates ($\dot{\gamma} = 20/\text{s}$).

We use periodic boundaries in x- and y-direction and closed boundaries in z-direction [17]. Shear is applied in x-direction by moving small zones of

particles and fluid close to the wall with a given shear velocity. The xy-plane is our shear plane. For simulations without shear, to achieve the best comparability, we use the same boundary conditions and just set the shear rate to $\dot{\gamma} = 0$. In addition we have performed simulations with two different shear rates: with $\dot{\gamma} = 100/\mathrm{s}$ and with $\dot{\gamma} = 500/\mathrm{s}$.

5 Results and Discussion

First, we focus on simulations without shear, where one can predict intuitively, what should happen. Qualitatively the results are similar to our earlier work [20], but the quantitative relation between the pH-value and the potentials is new. The relation was presented in [17], but here we apply it to different cases and we focus more on the characterization of the microstructure. However, given the particle particle interaction potentials, the microstructure in equilibrium can be predicted easily, at least on a qualitative level. But, the matter changes and gets more sophisticated, when shear is applied and an interplay between shear flow and particle particle interactions becomes responsible for the resulting microstructure. At the end of this section we move on to dilute suspensions and study the growth of clusters at low shear rates.

5.1 Correlation Function

For constant ionic strength $I = 3\,\mathrm{mmol/l}$ the local microstructure can be examined using the correlation function. Depending on the pH-value the behavior of the system changes from a repulsive structure around $p\mathrm{H} = 4$ to a stable suspension around $p\mathrm{H} = 6$ towards a clustered region if the pH-value is further increased, until the isoelectric point is reached at $p\mathrm{H} = 8.7$. There clustering occurs in any case, independent on the ionic strength. This can be seen in the structure of the correlation function

$$g(r) = \frac{V}{N^2} \left\langle \sum_i \sum_{j \neq i} \delta(r - r_{ij}) \right\rangle, \tag{7}$$

(see [5] p. 55), where V is the volume, N the number of particles and r_{ij} the distance of two particles i and j.

At $p\mathrm{H} = 4$ electrostatic repulsion prevents clustering: Particles are suspended, and there is no fixed long range ordering in the system. The correlation function (Fig. 1) shows a maximum at a typical nearest neighbor distance slightly above $\frac{r}{d} = 1$ with d denoting the particle diameter, then in the layer of next neighbors small correlations can be found (at $\frac{r}{d} = 2$). For larger distances the correlation function is rather constant.

When the pH-value is increased, the surface charge is lower, which at first causes the particles to approach each other more closely. The maximum of the correlation function is shifted to smaller distances (see Fig. 1, note that the

curves are shifted vertically in the plot by a factor of 3 for better visibility.). Then, van der Waals attraction becomes more important and clustering begins. One can see this in the correlation function where a sharp structure at particle distances between 1.5 and 2 particle diameters occurs. There is a nearest neighbor peak, and more complicated structures at larger distances, which we have assigned to typical particle configurations for small distances [20]. In a solid like cluster the position of the next neighbor is fixed more sharply than in the suspension, consequently the nearest neighbor peak becomes sharper, and its height is increased. Close to the isoelectric point (pH = 8.7) the barrier between primary and secondary minimum disappears. The particles, once clustered, cannot rearrange anymore, and therefore the correlations to the next neighbors become less sharp again (compare the cases of pH = 8.7 and pH = 7.7 in Fig. 1 at the positions denoted by the arrows).

Instead of varying the pH-value, one can also vary the ionic strength to achieve similar effects. Increasing the ionic strength, experimentally speaking "adding salt" decreases the screening length $1/\kappa$ and therefore the attractive forces become more important: the particles start to form clusters.

The effects described up to here can be observed with or without shear qualitatively in an analogous manner. If the suspension is sheared clustering occurs at higher pH-values and the peaks found in the correlation function are slightly broadened, because the relative particle positions are less fixed. But a new feature appears, if a stable suspension of not too high volume fraction is sheared. Induced by the shear particles arrange themselves in layers. Regular nearest neighbor distances in the shear plane cause the correlation function to

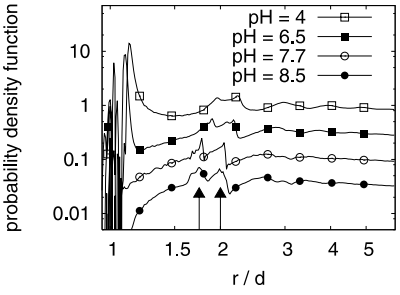

Fig. 1. Dependence of the particle correlation function on the pH value, $I = 3$ mmol, $\dot{\gamma} = 0/\mathrm{s}$ $\Phi = 35\%$. The plots for four different pH-values are shifted against each other for better visibility by a factor of 3. For pH = 4 the particles are not clustered. Hence the structure at $\frac{r}{d} = 2$ is less sharp than in the other three curves of the plot and the nearest neighbor peak (at $\frac{r}{d} = 1$) is broad. For pH = 6.5 slight clustering starts, the structures become sharper. For pH = 7.7 strong cluster formation is reflected in very sharp structures. For pH = 8.5 electrostatic repulsion nearly disappears so that no barrier between primary and secondary minimum exists anymore. The particles cannot rearrange anymore, and therefore the structures labeled by the arrows become smoothened compared to the case of pH = 7.7 (source: [19])

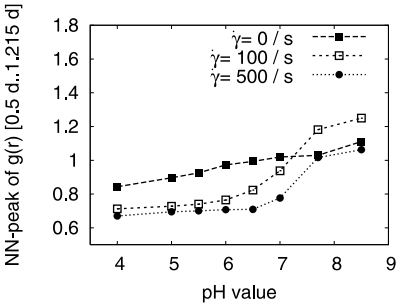

Fig. 2. Nearest neighbor peak (primary and secondary minimum of the potential) of the correlation function $I = 3\,\mathrm{mmol/l}$, $\Phi = 35\,\%$: For low pH-values clustering is prevented by the electrostatic repulsion. For high pH-values the particles form clusters, which is reflected by an increased nearest neighbor peak. First, shear prevents clustering, then depending on the shear rate, cluster formation takes place. Low shear rates even support cluster formation at high pH-values (source: [19])

become more structured even for large distances. The long range structure of the pair correlation function appears after a transient time the particles need to arrange themselves in the layered structure. Shear induced layer formation has been found in both, experiments [1, 2] and simulations [8, 9, 36]

We have integrated over the nearest neighbor peaks, both, the peaks of the primary and the secondary minimum, and plotted the integral versus pH-value in Fig. 2. We have chosen $I = 3\,\mathrm{mmol/l}$ and $\Phi = 35\,\%$ and three different shear rates: $\dot{\gamma} = 0, 100$ and $500/\mathrm{s}$. We have integrated the correlation function for $r < 1.215\,d$, where for all pH-values the potential in the secondary minimum has a value of $-\frac{1}{2}k_\mathrm{B}T$. In other words, we have captured the primary and the secondary minimum of the potential for this plot. For low pH-values clustering (in the secondary minimum) is only possible for low shear rates. For high shear rates, the hydrodynamic forces do not allow the formation of stable clusters. For rising pH-values the clustering increases, first for the un-sheared suspension, at higher pH-values for low shear rates ($\dot{\gamma} = 100/\mathrm{s}$) and finally for high shear rates ($\dot{\gamma} = 500/\mathrm{s}$). Remarkably, for $p\mathrm{H} > 7.5$ the curve for $\dot{\gamma} = 100/\mathrm{s}$ shows stronger cluster formation than the other ones. Particles are brought together by the shear flow, so that compared to the case of no shear, the clustering process is supported here. On the other hand, the shear stress may not be too strong, because otherwise the clustering process is limited by the shear flow again (for $\dot{\gamma} = 500/\mathrm{s}$ the clustering is less pronounced than for $\dot{\gamma} = 100/\mathrm{s}$).

5.2 Structure Factor

The pair correlation function can be used to characterize the local order of the microstructure on the length scale of the particle size. However, to do the characterization on the length scale of the system size, we use the structure

factor defined by

$$S(\mathbf{k}) = \frac{1}{N} \sum_{l,m=1}^{N} \exp(i\mathbf{k} \cdot \mathbf{r}_{lm}), \tag{8}$$

where N is the number of particles, and $\mathbf{r}_{lm}$ is the vector from particle l to particle m. i denotes the imaginary unit here. The structure factor is related to the pair correlation function in real space by a three dimensional Fourier transform. In principle the structure factor contains the same information as the pair correlation function. However, due to numerical reasons and our implementation of shear boundary conditions it is easier to observe the long-range structure in the structure factor than in the pair correlation function.

In Fig. 3 we have plotted several typical structure factors of our simulations. For these plots the pH-value is fixed to $p\mathrm{H} = 6$. The cases a) and b) are sheared with $\dot{\gamma} = 500/\mathrm{s}$ at an ionic strength of $I = 0.3\,\mathrm{mmol/l}$. In case a) the volume fraction $\Phi = 20\%$ is relatively low. Therefore the particles can arrange themselves in layers parallel to the shear plane, which move relatively independently in the shear flow. They have a certain distance fixed in space and time. This can be seen in a sharp peak at a dimensionless $\mathbf{k}$-vector of $k = 5.2$, which corresponds to a distance of 1.2 particle diameters. In fact, this is exactly the distance between two neighboring layers, as one can easily verify by counting the layers in a snapshot of the system (Fig. 4a)). The particles in the layers do not have a fixed distance and therefore no 2nd-order-peak can be observed.

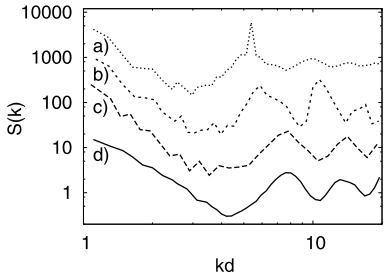

Fig. 3. Structure factor for some selected examples, with $p\mathrm{H} = 6$ fixed for all plots: $\dot{\gamma} = 500/\mathrm{s}$, $I = 0.3\,\mathrm{mmol/l}$: *a*) $\Phi = 20\%$ and *b*) $\Phi = 35\%$, $\dot{\gamma} = 0$, $I = 25\,\mathrm{mmol/l}$: *c*) $\Phi = 40\%$ and *d*) $\Phi = 10\%$. The curves are shifted vertically for better visibility. In case *a*) ten layers can be identified in the system, resulting in the strong peak close to 5. But, since the particles in the layers can still move freely, there is no 2nd-order-peak. In case *b*) layers are formed, but particles are moving from one layer to the other, disturbing the flow. As a result the nearest neighbor peak is much broader. Due to the structure in the layers, a 2nd-order-peak appears. In case *c*) the interaction is strongly attractive, hence the particles approach each other and the nearest neighbor peak is shifted to higher $\mathbf{k}$-vectors. In case *d*) the volume fraction is much less. The slope of the low-$\mathbf{k}$-peak is much flatter, which depicts that the cluster is fractal (source: [19])

For case b) the volume fraction is increased to $\Phi = 35\%$. The particle layers are packed more densely and therefore the interactions between one layer and the neighboring one become relevant. Particles jump from one layer to the other, which disturbs the flow and therefore the distance between the layers is not fixed anymore. The sharp peak on top of the nearest neighbor peak disappears. Instead of that, in each layer a regular hexagonal order appears and therefore the 2nd-order-peak is much more pronounced.

In case c) the ionic strength is increased to $I = 25\,\mathrm{mmol/l}$. The interparticle potentials are attractive enough that aggregation takes place. In this simulation we did not apply shear, therefore one finds only one big cluster in the system (compare Fig. 4c)). In the cluster the particles are packed more densely and consistently the nearest neighbor peak in the structure factor is shifted to larger **k**-vectors. The volume fraction is $\Phi = 40\%$ in this case.

In case d) the volume fraction is decreased to $\Phi = 10\%$. The particles still form clusters, but their mobility is not high enough to create one compact cluster. The system has a fractal structure (see Fig. 4d)). This can be seen in the structure factor as well: The slope for low **k**-vectors is flatter in this case compared to cases a)–c). A flatter slope of the low-**k**-peak is typical for structure factors of fractal objects. The fractal dimension of the cluster extracted from the slope of the low-**k**-peak is 2.5. In experiments this relation is often used to determine the fractal dimension of a sample: Lattuada et al. [31] have evaluated the fractal dimension of agglomerates of latex particles from the slope of the structure factor. McCarthy et al. [35] give an introduction to scattering intensities at fractal objects, without mentioning the structure factor, but their arguments refer to the contribution of the structure factor on the scattering intensity. The underlying mechanism which is responsible for these structures is cluster cluster aggregation [48].

In Fig. 5 we show the dependence of the low-**k**-peak of the structure factor on the *p*H-value. Here we have integrated over dimensionless **k**-vectors smaller than 3 which means, we have captured structures larger than twice a particle diameter. A large integral over the low-**k**-peak is due to a large inhomogeneity in the system. In one part of the system particles are present and in the other part not. In other words, we observe the process of cluster formation on a length scale of the system size. Without shear, particles cluster in the secondary minimum for all *p*H-values. If the system is slightly sheared ($\dot{\gamma} = 100/\mathrm{s}$) clustering is suppressed for low *p*H-values. Starting at $p\mathrm{H} = 6$ cluster formation starts and is even supported by the shear flow for *p*H-values larger than 7.5. For large shear rates ($\dot{\gamma} = 500/\mathrm{s}$) cluster formation is suppressed by the shear flow. By analyzing the low-**k**-peak of the structure factor one observes on the length scale of the system size. The same behavior of the system can be seen by analyzing the pair correlation function, as we have already shown in Fig. 2. In that case one analyzes the number of nearest neighbors, that means, one observes the length scale of a particle diameter. Nevertheless, both graphs show the same behavior of the system, i.e., we have

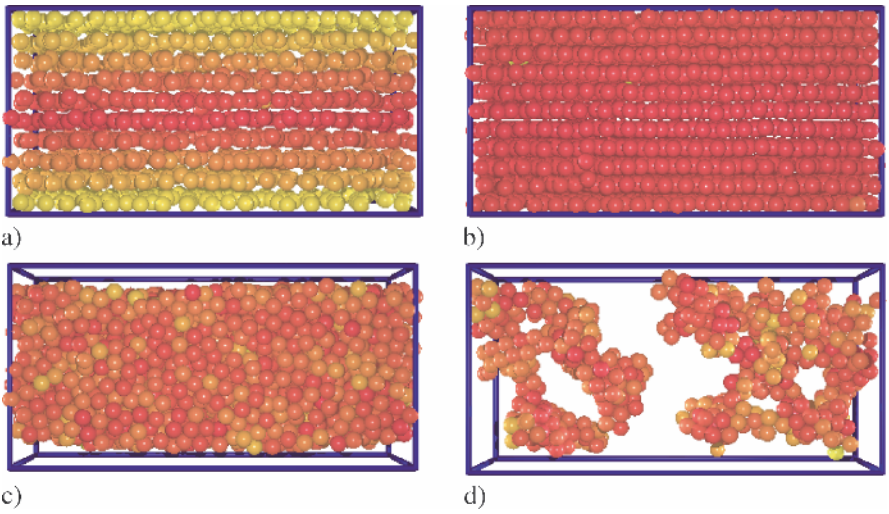

Fig. 4. Snapshots of the systems analyzed in Fig. 3: In case **a**) one can see the layers resulting in the sharp peak in the structure factor. In case **b**) the layers are packed closer due to the higher volume fraction. Collisions between particles of neighboring layers happen more frequently. In case **c**) one big cluster is formed. The particles are packed densely. In case **d**) the fractal nature of the system can be seen directly (source: [19])

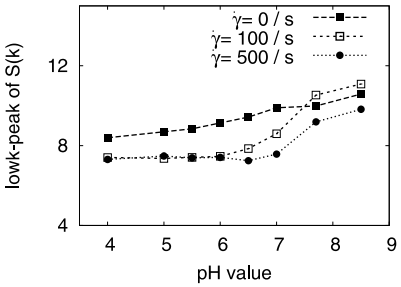

Fig. 5. low-**k**-peak for different *p*H-values and different shear rates. The ionic strength I is kept constant at $I = 3\,\mathrm{mmol/l}$ and the volume fraction is always $\Phi = 35\%$. For $\dot{\gamma} = 0/\mathrm{s}$ the particles tend to cluster in the secondary minimum of the potential. This clustering can easily be broken up, if shear is applied. If the *p*H-value is increased, shear cannot prevent cluster formation anymore. At low shear rates ($\dot{\gamma} = 100/\mathrm{s}$) clustering is even enhanced, since the particles are brought closer to each other by the shear flow (source: [19])