

*R. Backofen, H.-G. Borrmann, W. Deck, A. Dedner, L. De Raedt, K. Desch, M. Diesmann, M. Geier, A. Greiner, W.R. Hess, J. Honerkamp, St. Jankowski, I. Krossing, A.W. Liehr, A. Karwath, R. Klöfkor, R. Pesché, T. Potjans, M.C. Röttger, L. Schmidt-Thieme, G. Schneider, B. Voß, B. Wiebelt, P. Wienemann und V.-H. Winterer*¹

A Bottom-up approach to Grid-Computing at a University: the Black-Forest-Grid Initiative

ABSTRACT

Recent years have seen a rapid increase in the need for high-performance computing. These demands come from disciplines such as particle physics traditionally relying on High Performance Computing (HPC) but lately also from the various branches of life science that have matured into quantitative disciplines. The classical infrastructure of university computer centres results to be unsuited to cope with the new requirements for a multitude of reasons. Here we discuss the causes of this failure and present a solution developed at the University of Freiburg in a collaborative effort of several faculties. We demonstrate that using state of the art grid computing technology the problem can now be addressed in a bottom-up approach. The organizational, technical, and financial components of our framework, the Black Forest Grid Initiative (BFG) are described and results of its implementation are presented. In the process, a number of new questions have emerged which the next phase of our project needs to address.

MOTIVATION

The need for scientific high-performance computing (HPC) is ever-growing. The possibility of solving complex mathematical models by numerical simulation methods allows the study of physical, chemical, and biological systems at a level of detail otherwise not accessible. In many fields of science ranging from fundamental microscopic physics to brain research, the availability of sufficient computing resources is often the limiting factor for scientific progress. The list of the numerous large scale scientific projects includes among others the Human Proteome Folding, the United Devices Cancer Research Project, both hosted by the non-profit organization grid.org and the LHC (Large Hadron Collider) Computing Grid (LCG) [1]. Depending on the specific problems, the requirements imposed on the computing infrastructure differ significantly. One class of problems can be traced back to solving high-dimensional linear equation systems. These problems are highly parallel and can best be solved on dedicated parallel computing hardware with a high-speed low-latency inter-process communication. On the other hand, a huge class of problems is either completely serial in its nature or can be serialized without significant loss in efficiency and scalability. While the first class is currently typically attacked at HPC parallel computer centres such as HLRS Stuttgart etc., the latter class is traditionally solved on local computer clusters often operated by individual university research groups. Unfortunately, such local clusters often require a significant fraction of the research group's personnel resources. Furthermore, cluster solutions are often not easily scalable.

Commodity hardware offers in principle the possibility to accumulate both very high CPU power and huge storage resources at affordable prices. However, reliability and operability of clusters in excess of 1000 CPUs and tens to hundreds of Terabyte of disk space impose serious constraints which have to be solved in a well-thought and professional manner.

Computing resources managed by individual research groups are seldom used efficiently, since the demand varies significantly with time. Sharing of resources between different research groups within a university and sharing of resources between different collaborating research institutions can increase the available peak computing power dramatically. This resource sharing is the basic idea of grid computing [2], [3]. While ideally, grid computing would provide resources to the user in a trans-

¹ *Rolf Backofen, Luc De Raedt, Stefan Jankowski, Andreas Karwath, Lars Schmidt-Thieme*: Institut für Informatik, Universität Freiburg, Georges-Koehler-Allee 51, 79110 Freiburg, Germany.

Hans-Gunther Borrmann, Raphaël Pesché, G. Schneider, Volker-Henning Winterer: Rechenzentrum der Universität Freiburg, Hermann-Herder-Str. 10, 79104 Freiburg, Germany.

Werner Deck, Ingo Krossing: Institut für Anorganische und Analytische Chemie, Universität Freiburg, Albertstr. 21, 79104 Freiburg, Germany.

Klaus Desch, Tobias Potjans, Peter Wienemann: Physikalisches Institut, Universität Freiburg, Hermann-Herder-Str. 3, 79104 Freiburg, Germany.

Markus Diesmann, Bernd Wiebelt: Bernstein Center for Computational Neuroscience, Universität Freiburg, Hansastr. 9a, 79104 Freiburg, Germany.

Martin Geier, Andreas Greiner: Institut für Mikrosystemtechnik, Universität Freiburg, Georges-Koehler-Allee 51, 79110 Freiburg, Germany.

Wolfgang R. Hess, Björn Voß: Institut für Biologie II, Universität Freiburg, Schänzlestr. 1, 79104 Freiburg, Germany.

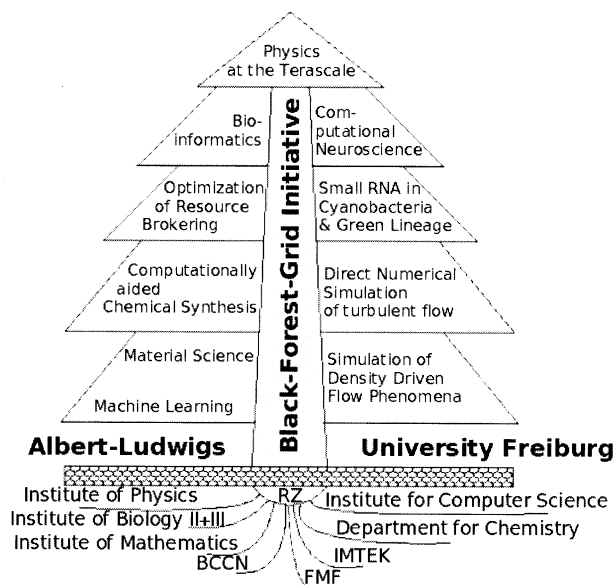
Josef Honerkamp, Andreas W. Liehr, Michael C. Röttger: Freiburger Materialforschungszentrum, Universität Freiburg, Stefan-Meier-Str. 21, 79104 Freiburg, Germany.

Andreas Dedner, Robert Klöfkor: Mathematisches Institut, Universität Freiburg, Abteilung für Angewandte Mathematik, Hermann-Herder-Str. 10, 79104 Freiburg, Germany.

Markus Diesmann: Institut für Biologie III, Universität Freiburg, Schänzlestr. 1, 79104 Freiburg, Germany.

parent and platform-independent way, it appears to be difficult to accommodate the different, often platform-dependent needs of the users.

We take a pragmatic bottom-up approach to a computing grid in our Black Forest Grid Initiative (BFG) at the University of Freiburg. It is based on basic agreements on hardware and operating system by the participating research groups without making the attempt to arrive at a minimal consensus which could accommodate *all* needs. Rather we impose the most stringent constraints which can still fulfil the largest fraction of the requirements. Furthermore, we rely on the grid paradigm from the beginning. Therefore users will not be able to use the resources interactively or via traditional local batch systems but only through the grid middleware. By choosing this approach we ensure that the installation is scalable and can be easily interfaced with future external resources from collaborators of the different participating research groups. Also user administration remains manageable in this way.



Organization scheme of the BFG Initiative: cooperating institutes and groups, areas of application.

CHOICES

Hardware

We have chosen standard dual-CPU rack-mounted compute servers connected by GBit-Ethernet as computing hardware. This choice was driven mainly by cost-effectiveness. The first batch of 25 nodes are SunFire V20z servers with two 2.2 GHz Opteron 248 processors equipped with 2 GB RAM and 70 GB SCSI hard disks. Recently the cluster has been extended by 34 SunFire X4100 machines containing two 2.2 GHz Opteron 248 CPUs, 4 GB RAM and 36 GB SAS hard disks. Disk storage is provided by 3.6 TB RAID6 external storage server equipped with twelve 400 GB SATA disks. The disk server is connected via SCSI to one of the compute servers. While further hardware upgrades do not have to be identical to this choice, there is the clear aim to minimize the heterogeneity of the system as much as possible.

Operating System, Cluster Management Software and Grid Middleware

According to the requirements of the vast majority of participating research groups we have chosen Linux as the operating system. Currently we are running the Scientific Linux CERN SLC3.0.5 distribution which is freely available from CERN [4]. This choice is motivated by its easy availability and known compatibility with the grid middleware.

In order to minimize the cluster administration effort we have chosen the quattor cluster management system [5] to setup and update the system. Quattor is used e.g. at CERN to successfully operate clusters of more than 2400 nodes. It allows combining autonomous nodes (local configuration files, no reliance on global file systems such as NFS or AFS) with a single source for all configuration information. The basis of quattor's information system is the distinction between the *desired* and the *actual* state of the nodes. The desired state is stored in a central configuration database (CDB) and from there it is propagated and cached on the managed nodes. Services running on each node take care of managing software packages and configuring local services.

The key choice of the installation is the grid middleware. Due to the growing efforts in Germany, Europe and world-wide, various systems are available for scientific grid computing (e.g. UNICORE [6], globus [7]). We rely on the LCG (LHC Computing Grid) middleware [1]. LCG is based on the tools developed in the EU-funded projects EDG (European Data Grid) and EGEE (Enabling Grids for EScience in Europe) and is currently deployed by the LCG deployment team at CERN [8]. The LCG installation is meanwhile the largest scientific grid installation world-wide consisting in total of 16.000 contributed CPUs and 4.000 TB of storage. One of its main applications will be the analysis of the data that will be provided by the Large Hadron Collider LHC, the flagship project in collider-based particle physics in the coming decade. This choice is partly motivated by the fact that one of the participating groups at the University of Freiburg is member of one of the LHC experiments, the ATLAS experiment, and thus has to rely on LCG. However it turns out that choosing LCG is also beneficial for the other research groups participating in the BFG initiative. The main reasons for this are the wide-spread use of LCG which results in both very good user support and rapid development. Finally, we benefit from the fact the LHC project will start operating in 2007 and thus a huge prototype application of significantly larger size than BFG can provide valuable operating experience.

ORGANIZATION

The BFG hardware is hosted by the Computing Center of the University of Freiburg, Germany. Hardware and software choices, models for sharing of the resources as well as integration of new research groups are discussed within open BFG meetings. Decisions are taken on the basis of mutual consensus among the partners. Funding of the initial configuration was shared by the research groups and the Computing Center. The operating model is similar to that of large scientific collaborations, e.g. in particle physics. With a growing number of partners, this model will be formalized in the near future without reducing the great flexibility of the present model.

It should be noted that this flat, 'shareholder'-like operation model maximizes the scientific usability of the system and minimizes any administrative overhead.

STATUS

In order to put the grid dream – transparent access to computing and storage resources spread around the world – into practice, a number of services must be provided by the middleware that shields the complexity from the user. This requires sophisticated information and authentication systems.

Every LCG site publishes the provided resources to the outside world through its site GII (Grid Index Information Server). The information from site GIISes is collected by BDIs (Berkeley Database Information Index) which can be interrogated to retrieve all sites providing certain wanted resources.

Computing resources are provided to the grid through computing elements (CE), storage space through storage elements (SE). To allow easy access to stored data distributed among many physical storage devices around the globe, a file catalogue (LFC, LCG file catalogue) is provided which maps easy-to-use logical file names to storage URLs.

Users access grid resources through tools made available on user interfaces (UI) and worker nodes (WN). Each user needs a cryptographic X.509 certificate from an LCG-approved certification authority (CA) to access grid resources. Prior to grid operations, users need to create a proxy certificate from the user certificate. A proxy certificate is a delegated user credential of limited lifetime which authenticates the user in every secure interaction. Submitted jobs are first moved to a resource broker (RB) which decides, based on information provided by the submitter, which of the available resources fits best the job requirements. After this match-making process, the RB finally chooses a CE which in turn submits the jobs to a local resource management system (LRMS).

The grid user community is divided into virtual organizations (VO). It is a compulsory requirement for all users to belong to a VO. Since suitable VOs were not available for all BFG member groups, new VOs hosted in Freiburg have been set up for computational neuroscience (cns) and, to combine users of a number of smaller groups, the VO BFG. In addition to these locally hosted VOs, the VOs dteam (LCG deployment team), atlas (ATLAS high-energy physics experiment), ilc (international linear collider project) are supported. All LCG services obtain their information required for user authentication from a VO membership service (VOMS) server.

The BFG site runs almost the full spectrum of LCG services and a dCache [9] storage element. This system provides the grid community with a powerful tool for storing and retrieving huge amounts of data. It allows to access data potentially distributed among heterogeneous server nodes under a single virtual file system tree and provides fail-safe storage space management as well as sophisticated load balancing.

The BFG site is a certified LCG site since November 15, 2005 and has provided 27000 kSPECint2000 hours during the first three months since certification according to the accounting information published by the Grid Operations Centre (GOC) [10].

APPLICATIONS

Bioinformatics

Lehrstuhl für Bioinformatik, Institut für Informatik

The bioinformatics research group deals with diverse problems of the field of bioinformatics. These include the prediction and

analysis of alternative spliceforms and the description and detection of regulatory elements.

Furthermore, the research group has developed several methods for detecting functional RNA motifs using sequence/structure alignment. The multiple sequence/structure alignment tool MARNAs [11] is currently one of the best available systems for this purpose. On a current single-CPU system, the MARNAs system allows the processing of a limited number of RNA sequences of up to approx. 500 nucleotides each. Use of the grid resources will allow MARNAs to scale up to comparing a higher number of much longer sequences, as is frequently requested. In addition, the protein folding problem is investigated on simplified models of proteins. Protein structure prediction, even in simplified models, is a NP-complete problem. Since structure and sequence space grow exponentially with the protein size, distributed computing is a necessity in order to handle proteins of realistic size. The inverse RNA folding problem is yet another field of interest of the group. A fast and successful RNA design tool was developed. Last but not least, the group is interested in selenoproteins and the design of the selenocysteine insertion element (SECIS).

In all of these fields, use of the increased computing power offered by the grid allows to tackle much larger and more complex problem instances.

Computational Neuroscience

Bernstein Center for Computational Neuroscience (BCCN)

Higher brain functions are held to arise through the cortical neuronal networks of the mammalian brain. This substrate is among the most complex found in nature. There are of the order of 100,000 nerve cells in a cubic millimetres of cortical tissue. With a probability of 0.1 of forming a local contact (synapse) between cells, the number of inputs a cell receives is already in the order of 10,000. The experimentally observed activity is characterized by phenomena over a wide range of spatial and temporal scales, from the micrometer scale of a single nerve cell to the centimetre scale of cortical areas, and from the submillisecond scale of events at individual synapses to the scale of minutes at which learning occurs. It is part of the BCCN's mission to develop the theoretical framework and computational techniques required to understand the dynamical and functional properties of these networks [12]. The BCCN also hosts the software infrastructure of the NEST initiative [13] dedicated to the development of simulation technology for large scale neuronal networks.

The need to perform large scale simulations and to identify the development of the corresponding numerical techniques as a scientific problem in its own right has long been underestimated (but see [14], [15]). However, several groups, specifically in Europe, are now engaged in this endeavour. Next to large scale simulations of neuronal networks, Computational Neuroscience relies on the analysis of multi-dimensional experimental data characterized by non-stationarity in time and across trials. Monte Carlo methods and advanced statistical techniques like time resolved correlation analysis are required to interpret the data. Often the same analysis needs to be independently carried out with different parameter settings on a large number of data sets. Therefore, the BFG represents the ideal computing infrastructure for the applications of the BCCN. The BCCN Freiburg is one component in a national network of Centers for Computational Neuroscience [16]. It needs to be explored

whether in the frame work of the VO cns computing resources of the computational neuroscience community can be shared on a national and international level.

Computationally aided chemical syntheses

Institut für Anorganische und Analytische Chemie

Is it possible to synthesize a yet unknown chemical compound? And if so, does it have the desired properties? These questions may be answered in the laboratory through tedious, expensive and time-consuming experiments. Alternatively such questions may also be answered with the help of quantum chemical calculations, even before the syntheses are carried out. Thus, computational chemistry is a valuable tool, helping us in the development of new weakly coordinating anions and tailor their properties [17] for diverse applications such as homogenous catalysis, Li ion batteries, Ionic Liquids but also the stabilisation of very reactive cations and weak Lewis acid-base adducts. Due to the involved weak and dispersive interactions, the latter compounds are often hard to model computationally [18]. Therefore enough computational power for sophisticated computational methods is needed to accurately study the properties of such compounds.

Participation in the Black Forest Grid project gives our group access to computational resources which enable us to calculate larger molecules than ever or to assess problem cases for theory with the highest achievable accuracy.

Direct Numerical Simulation of turbulent flow using Digital Lattice Boltzmann Automata


Institut für Mikrosystemtechnik

Turbulence is the chaotic behaviour of large fluidic systems induced by complex interactions with external forces. Due to its non-linear nature depending on huge amounts of variables, it is one of the least understood phenomena in nature. Current fluid dynamic simulation tools try to circumvent turbulence simulation by introducing phenomenological turbulence models. However, Direct Simulation of turbulence is the only way to gain new insight into the phenomenon itself. Our classical understanding of physical phenomena requires them to be investigated by reproducible experiments. We try to describe systems with well-defined measurable variables. These variables are supposed to settle to constant or oscillating values after a short time. For turbulence, nothing of the above is true. Turbulence is by definition irreproducible, depends on the size of the system and it does not converge to oscillating stages with periodicity shorter than the lifetime of the universe. Furthermore, it cannot easily be separated from outside effects such as the gravity of the Andromeda strain. Direct Numerical Simulation is in fact the only option to investigate turbulence in a reproducible and isolated way. Direct Numerical Simulation of turbulent flow is extraordinary expensive in terms of the required computational power. At least tens to hundreds of millions of degrees of freedom are necessary to capture the dynamics of three dimensional turbulence. This becomes slowly accomplishable by recent progress in both fields of computation, software [19] and hardware.

The Cellular Automaton based algorithm called the Cascaded Digital Lattice Boltzmann Method, which was recently developed by the Imtek simulation group [20], has superior stability characteristics as compared to other fluidic simulation methods without compromising physical soundness. It employs

Cartesian grids to fit several hundreds of millions of degrees of freedom into the main memory of a standard workstation. The method applies exact arithmetic and is hence free of round-off errors. Simulations can run for an infinite number of time steps without becoming unstable.

Turbulence is irreproducible. Isolated simulations provide little insight to the subject. Averaging techniques are necessary to investigate the statistical behaviour of turbulent flow. Due to limited computing resources only spatial and temporal averaging has been applied so far. With these techniques spatial and temporal resolution is lost. To obtain detailed information on the “where” and “when” of turbulent bifurcations and to analyze the predictability of a given flow situation ensemble averages are necessary. Grid computing provides the necessary computational resources. It allows us to run the same simulation with slightly different initial conditions many times in parallel to obtain good statistical data from which we can gain new insights in the onset and evolution of turbulence in a time and space resolved way.



Simulation of turbulent fluid flow requires exceptional computational resources due to the complexity of the problem. This picture shows the turbulence in the wake behind an obstacle simulated with the Digital Lattice Boltzmann Automaton.

Machine Learning

Computer-based New Media group (CGNM),
Institut für Informatik

The Computer-based New Media group (CGNM) of the department of computer science uses the grid infrastructure for machine learning, especially for the computation of complex ensemble based classifiers with millions of component models. Such models are used, for example, for recommender systems in e-commerce which compute customer-specific preferences for thousands of different products [21].

Recommender systems are a personalization technique that adapts information systems to individual customers. Based on preference indicators for specific products or types of products, as past purchases, catalogue searches, viewing of products etc., a recommender system tries to predict which other products customers might be interested in and recommends them for buying. Although traditionally simple nearest neighbour models called collaborative filtering are used for these applications, more powerful classifiers such as support vector machines can be used.

From a computational point of view the main problem is the large number of alternative classes that the classifier has to discriminate. While learning multi-category classifiers with 3 to 26 classes, e.g., to discriminate the different letters in the alphabet for handwriting recognition, is standard, classifiers for hundreds or thousands of classes are computationally demanding as there is a runtime component that is quadratic in the number of

classes. On the other hand multi-category problems are trivially parallelizable as standard model setups, e.g., one-vs.-rest or one-vs.-one structure the problem in an ensemble of binary subproblems. So Grid technologies gracefully help to compute solutions in time.

Material Science

Service-Group Scientific Data Processing,
Freiburger Materialforschungszentrum (FMF)

If one looks at an unknown material the most important property being recognized is the color of the reflected light. This intuitive material characterization can be utilized by so-called reflectance spectrometry. Here the material is irradiated with a white light source of known intensity spectrum and the intensity of the reflected light is measured for each wavelength. This spectral reflectance reveals a lot of subtle information about the microstructure of the material like the concentration of light absorbing substances or the size distribution of light scattering obstacles. In order to gain this not directly accessible information one has to relate the mesoscopic material properties to the optical properties of the material and furthermore model the dependency of the reflectance from the optical material properties. The latter are parameterized by the scattering coefficient, the absorption coefficient and the anisotropy factor. While the absorption coefficient is determined by the concentration and the extinction spectrum of the light absorbing substance, the scattering coefficient and the anisotropy factor can be modelled in terms of the Mie theory if one assumes the shape of the light scattering obstacles to be spherical. Now the dependency of the reflectance on the optical material parameters can be determined by a Monte-Carlo model simulating the light transport in turbid media. Given this correlation the microscopic parameters can be estimated from a measured reflectance spectrum by least-square or regularization methods [22], [23]. In order to analyse the measured spectral reflectance of a certain type of material up to 1000 Monte-Carlo simulations have to be carried out with each simulation covering one set of optical parameters. Therefore one deals with a classical parameter scan scenario which can be distributed on a grid computing environment like the BFG. Currently these computing tasks are distributed by a local resource management system to the workstations of the Freiburg Materialforschungszentrum. The connection of these resources to the BFG via a Condor grid [24], [25] is scheduled.

Optimization of Resource Brokering

Service-Group Scientific Data Processing,
Freiburger Materialforschungszentrum (FMF)

The construction of the Grid poses many challenges in fields of organisation and software engineering. One of the open questions is how to distribute jobs to resources optimally. This mapping is done by the so-called "broker". The question is, how to decide which strategy is useful and how can it be found? Should it be completely decentral? Is it possible to implement a strategy, which is able to adjust itself to a constantly changing computing environment? It is not promising to tackle all these problems at once, so we're starting with an elementary model of a network with quite simple resources and outstanding jobs distributed by one central broker.

Our first approach starts with a simplified view of the workload management system [25] of the University of Wisconsin/Madi-

son, Condor's brokering unit. This so-called *Matchmaker* periodically gets information about available resources with their capability characteristics, and about outstanding jobs in the queues with their requirements on computational power and storage. Among all possible matches, considering performances and demands, the matchmaker chooses pairs of resources and jobs and briefs the involved computers and queues about its decision. The latter arrange themselves in order to get the jobs done.

To find an optimal strategy for the matchmaker in our model, we are using ideas and methods of *Reinforcement Learning* (RL) (e.g. [26]), which implies a special view to the system: One agent called *Negotiator* is periodically informed about the state of his environment, the grid. After every action, the agent receives a reward, which gives him feedback about the utility of his doings. An efficient learning algorithm has to cope with e.g. the generalization of the continuous states of the system and the varying number and types of actions among which the agent can choose. In the course of time, the agent tries to choose actions in order to maximize the expected sum of his rewards. Our goal is to implement the appropriate rewards and learning algorithms in order to find a suitable strategy, which can be applied to a real subset of the grid. The Black Forest Grid is such a subset in addition to a local Condor pool.

Simulation and Analysis of Particle Collisions at the Terascale

Physikalisches Institut

The aim of particle physics is to explore the fundamental nature of matter and the basic forces that shape our universe.

In the years ahead of us both theory and experiment suggest that new groundbreaking discoveries will come in an energy range called the Terascale, an energy that an electron gains when accelerated by a voltage of 1.000.000.000.000 volts (1 TeV). At the Terascale, new forces come into play which may explain the microscopic origin of mass, which may open a way to understand the nature of the astrophysically observed dark matter in our universe and may even show a path towards the ultimate unification of all forces.

Tremendous experimental efforts are required to access the Terascale. Particle accelerators play a special role in this enterprise. The Large Hadron Collider LHC [27] will turn into operation in the year 2007. The detectors ATLAS and CMS will be able to analyse and record 40 million proton-proton collisions per second at 14 TeV of centre-of-mass energy. Unfortunately, collisions which contain the decay products of new particles such as the Higgs Boson or supersymmetric particles are suppressed by factors up to 10^{10} . While a huge fraction of the collisions can be discarded online, a significant amount still has to be stored on disk and compared to simulated collisions. These experimental data make up for a huge data volume (10 TB/day) and also call for enormous computing resources to analyse them.

Within the LHC experiments the only solution to this is to pursue the GRID paradigm of globally distributed resources linked together via a transparent middleware, LCG. The Black Forest Grid installation in Freiburg is fully compliant with LCG and is a certified LCG site since November 2005. Through LCG, it is possible to share the BFG resources with other colleagues around the world within the ATLAS virtual organisation (VO) and in turn gain access to their resources. Within the coming

years, the computing infrastructure for the ATLAS VO in Freiburg is planned to be ramped up significantly in order to serve as a Tier-2 centre for the ATLAS experiment. Partners in Germany in this enterprise are DESY, Universität Wuppertal, LMU München, and Max-Planck-Institut für Physik, München. Furthermore, the high-energy physics group is involved in the R&D for a new electron positron linear collider, the ILC [28]. Within the ILC community, huge amounts of simulated data have to be shared in order to optimize the detector design.

Furthermore the particle physics group in Freiburg is involved in the D-GRID initiative [29]. Within these activities the data management tools required in the near future for the LHC experiments to store and provide petabytes of data are developed and tested. This includes the development of a scalable and tool-flexible test suite for dCache storage elements. The tests carried out give feed-back to the developers and help improving the grid infrastructure.

Simulation of Density Driven Flow Phenomena

Abteilung für Angewandte Mathematik

The need for simulating density driven flows arise in many fields ranging from large scale problems in astrophysics to microscale problems in industrial applications. Fast but accurate solvers are thereby essential for determining important physical parameters. This may require hundreds of simulations and a high amount of data which has to be stored. Also each run may require many time steps and a high spatial resolution – at least near shock formations – leading to hundreds of millions of degrees of freedom; here parallelisation is a central tool indispensable for high resolution and efficient simulations.

We are developing a general framework for solving this type of problem using the local Discontinuous Galerkin approach. Although a wide range of applications has to be covered, a common ground for the underlying mathematical models can be found. This observation suggests to define a common interface for the numerical schemes and to provide a generic implementation for central parts of the discretization. The scheme is implemented in the unified grid interface package DUNE, which includes both structured and unstructured grids in one, two, and three space dimensions.

One of the major goals of such an interface based numerics environment is the separation of data structures and algorithms. For instance, the problem implementation can be done on the basis of the interface independent of the data structure that is used for a specific application. Moreover, such a concept allows a reuse of existing codes beyond the interface. Also the required overhead for parallelization can be easily realized through the interface requiring little too no adaptation of the code to the specific application. This general framework is used in many projects at the Institute for Applied Mathematics with industrial and physical applications [30].

This approach enables us to develop solvers for solving complex mathematical models, for example a complex model describing the dynamic behaviour of a fuel cell stack. To perform parameter studies for a better understanding of these difficult models and the ongoing physical processes within a fuel cell, obviously a large amount of computer resources is needed. Due to the idea of sharing resources with other colleagues from all over the world, Grid Computing in general and the Black Forest Grid in particular will provide the necessary resources to

perform a large number of serial and parallel simulation runs using different parameter sets.

Small RNAs in Cyanobacteria and the Green Lineage

Experimental Bioinformatics, Institute of Biology II

Several total genome analysis projects have shown, that for many species including humans the number of protein-coding genes is far less than expected. Additionally, the protein coding fraction of genomes is often small (~5% in human), raising the question: "What is the remainder good for?". One answer is that the non-protein-coding fraction of the genome is important and that small RNAs encoded therein have key regulatory/catalytic functions. For a long time RNA (ribonucleic acid) was only regarded as a passive transporter of genetic information (messenger RNA) or as modules for protein synthesis (transfer RNA and ribosomal RNA). This view changed with the discovery of catalytic active RNAs, so-called ribozymes. Very recently, it was discovered that RNA can also have regulatory function. This class of RNA molecules is termed "ncRNAs" for non-coding RNAs or simply "small RNAs" or sRNAs for short.

This is true for higher organisms as well as for their evolutionary ancestors: For plants one such ancestral group are Cyanobacteria (photosynthetic bacteria; sometimes called blue-green algae) which have tremendously contributed to the higher plant gene pool. Today, four plant and about forty cyanobacterial genomes have been sequenced, which allow comparative studies to identify novel ncRNAs. As most biologically active compounds, also ncRNAs need a specific structure for correct function. This structure is encoded in the sequence and can be predicted. From an evolutionary perspective the problem arises, that two structurally identical RNAs do not need to have identical sequence, e.g. all tRNAs look like a cloverleaf but have different sequences. Thus, for comparative sequence analysis it is not sufficient to look for sequence conservation, one also has to take structure into account. Structure prediction is computationally more expensive (cubic time complexity) than sequence comparison (max. quadratic time complexity), resulting in the need for large computational resources, such as they are available with the BFG Initiative [31], [32].

Data Mining

Lehrstuhl für maschinelles Lernen und natürlichsprachlicher Systeme

The Lab for Machine Learning is active in various research areas, which all are computationally intensive. These include: the development of novel data mining methods that can be applied to challenging problems in bio- and chemo-informatics, network analysis, discovery, context-awareness, and robotics. Data mining techniques allow one to extract novel and useful patterns or knowledge through the analysis of large and possibly complex datasets.

Research in the lab focuses on two key research themes, which all aim at dealing with structured data that arise in the above application domains. The first theme is concerned with statistical relational learning, which combines expressive knowledge representation formalisms that use both relational logic and probabilistic models with techniques for learning these representations [33]. It is also the subject of the EU IST FET project APRIL II (Application of Probabilistic Inductive Logic Programming), coordinated by the lab. Various statistical relational mod-

els that have been developed include: logical hidden Markov models, Bayesian logic programs, and techniques for relational reinforcement learning.

The second theme is concerned with the development of powerful inductive query languages inductive queries are to be contrasted with the typical kind of queries one finds in a relational database, which merely retrieve information that resides in the database. Inductive queries, on the other hand, allow one to specify the patterns or models to be extracted from data in a declarative manner [34]. Using an inductive query, one can for instance ask for all chemical fragments that are present in at least 15 per cent of the active molecules and at most 1 per cent of the inactive ones.

To enable the application of statistical relational learning and inductive querying on scientific applications the lab employs the powerful facilities of the Black Forest Grid.

CONCLUSIONS

A novel organization form for the provisioning of the ever-growing computing power demand at universities has been realized. In contrast to classical top-down structures in HPC involving steering committees and rather rigid regulations in the operation of HPC centres, we adopted an adaptive and collaborative strategy to setup and run a HPC site. Instead of investing a high volume at a definite time, we rely on continuous relatively low to average volume investments. In doing this, we are able to efficiently take advantage of the rapidly evolving computer market. On the other hand, due to the heterogeneity of our collaborators and subsequently our wider contact range, we have an enhanced chance to foster collaborations in order to expand the Grid. We also expect an accentuated impact on the quality of education in respect to computational science at our university: due to the multiplicity of first-hand experiences in grid computing in the many collaborating groups of our initiative, an added value to the tutorials, courses and lectures widely spread over the disciplines can already be denoted. We actively contribute to that trend in organizing specific workshops and tutorials.

Furthermore, our Black Forest Grid Initiative provides feedback to science: as described in the applications chapter, several activities are under way to improve the quality of the Grid itself by investigating new strategies in the area of resource brokering and by optimizing other aspects of the Grid middleware. Various other research projects around the Grid, such as the study of serialization strategies for parallel algorithms and an investigation of the concept of using the grid as a grid of clusters for parallel algorithms with moderate communication load, are in a planning stage.

BCCN is partially funded by DIP F1.2 and BMBF Grant 01GQ0420 to the Bernstein Center for Computational Neuroscience Freiburg.

REFERENCES

- [1] <http://lcg.web.cern.ch/LCG/>
- [2] Roger Smith: Grid Computing: A Brief Technology Analysis, CTO Network Library, 2005.
- [3] Ian Foster; Car Kesselman: The Grid: Blueprint for a New Computing Infrastructure, Morgan Kaufmann publishers, 1999.
- [4] Scientific Linux Cern 3, <http://linux.web.cern.ch/linux/scientific3/>
- [5] Quattor system administration toolsuite, <http://quattor.org>
- [6] UNICORE, <http://www.unicore.org>
- [7] The Globus toolkit, <http://www.globus.org>
- [8] LCG Deployment, <http://lcg.web.cern.ch/LCG/activities/deployment.html>
- [9] dCache, <http://www.dcache.org>
- [10] LCG Grid Operations Centre, <http://goc.grid.sinica.edu.tw/goc/>
- [11] Siebert, S.; Backofen, R. (2005): MARNA: multiple alignment and consensus structure prediction of RNAs based on sequence structure comparisons. *Bioinformatics*, 21 (16), 3352-9.
- [12] www.bccn.uni-freiburg.de
- [13] www.nest-initiative.org
- [14] Morrison, A.; Mehring, C.; Geisel, T.; Aertsen, A.; Diesmann, M. (2005): Advancing the boundaries of high connectivity network simulation with distributed computing. *Neural Computing* 17 (8): 1776-1801.
- [15] Giles J.: Blue Brain boots up to mixed response. *Nature* 435 (2005), 720-721.
- [16] Bernstein Centers, <http://www.bernstein-centers.de/>
- [17] I. Crossing*; I. Raabe (2004): "A Computational Study of Weakly Coordinating Anions: Relative Stabilities, Acidities of their Parent Lewis Acids", *Chem. Eur. J.* 10, 5017-5031; H.D.B. Jenkins, I. Crossing*, J. Passmore, I. Raabe (2004): "On the Formation of Monomeric SbF₅ from Sb_nF_{5n} (n = 2, 3, 4). Implications for the Fluoride Ion Affinity of n SbF₅. A Computational Study*"; *J. Fluor. Chem.* 125, 1585.
- [18] M. Gonsior; I. Crossing* (2005): "A Theoretical Study on the Formation of EX₄⁺ and E₂X₅⁺ (E = P, As; X = Br, I) from Ag⁺ and EX₃/X₂." *Dalton Trans.* 2022-2030; H. Donald B. Jenkins, Luminita C. Jitariu, Ingo Crossing*, Jack Passmore, Reijo Suontamo (2000): "Basis Set and Correlation Effects in the Calculation of Accurate Gas Phase Dimerisation Energies of two M₂⁺ to give M₄²⁺", *J. Comp. Chem.*, 21, 218-226.
- [19] S. Ansumali; S.S. Chikatamarla; C.E. Frouzakis; K.B. Boulouchos: Entropic Lattice Boltzmann Simulation of the Flow Past Square Cylinder, *Int. J. Mod. Phys. C* (2004) 15, 435-445.
- [20] IMTEK simulation group, <http://www.imtek.de/simulation>
- [21] L. Schmidt-Thieme: Compound Classification Models for Recommender Systems. In: *Proceedings of IEEE International Conference on Data Mining (ICDM)*, Houston, USA (2005).
- [22] M. Bohnert; R. Walther; T. Roths; J. Honerkamp: *Int J Legal Med* 119, 355 (2005).
- [23] R. Walther; T. Roths; M. Bohnert; J. Honerkamp. In: *Monte Carlo based model for steady-state diffuse reflectance spectroscopy*, edited by SPIE Int. Soc. Opt. Eng (2004), 88.
- [24] D. Thain; M. Livny. In: *Building Reliable Clients and Services*, Edited by I. Foster and C. Kesselman (Elsevier, Amsterdam, 2004), 285.
- [25] Condor High Throughput Computing. <http://www.cs.wisc.edu/condor/> (2006).
- [26] Richard S. Sutton; Andrew G. Barto: *Reinforcement Learning. An Introduction.* (MIT Press, 1998).
- [27] The Large Hadron Collider Project, <http://lhc.web.cern.ch/lhc/>
- [28] The International Linear Collider, <http://www.linearcollider.org>
- [29] D-GRID initiative, <http://www.d-grid.de>
- [30] <http://www.mathematik.uni-freiburg.de/IAM/Research/projects/skr/dune/>
- [31] Axmann, I.M.; Kensche, P.; Vogel, J.; Kohl, S.; Herzel, H.; Hess, W.R. (2005): Identification of cyanobacterial non-coding RNAs by comparative genome analysis. *Genome Biology*, 6(9), R73.
- [32] Giegerich, R.; Voß, B.; Rehmsmeier, M.: Abstract shapes of RNA, *Nucleic Acids Research*, 32 (16), 4843-41 (2004).
- [33] L. De Raedt; K. Kersting (2003): Probabilistic Logic Learning. *SIGKDD Explorations* 5/1, pp. 31-48.
- [34] L. De Raedt (2001): A Perspective on Inductive Databases. *SIGKDD Explorations* 4/2, pp. 69-77.