WISSENSCHAFTSRAT

Recommendation on the Installation of European Supercomputers

Drs. 6341/04
Hamburg, 12 November 2004

Recommendation on the Installation of European Supercomputers

Summary

In its recommendation on the use of supercomputers in the future,[1)] the Wissenschaftsrat has pointed out that a continuous qualitative and quantitative upgrading of the computing infrastructure is essential for the safeguarding of scientific competitiveness. In accordance with the recommendation, the extension of supercomputing systems in Germany is happening with a replacement cycle of two to three years so that at least one system of the highest performance class is available at any time. This acquisition strategy has led to an adequate availability of supercomputer capacities in Germany. However, supercomputers of the highest performance class are necessary to compete with Japan, the USA and also China in the future.

Supercomputers of the highest performance class are indispensable for the numerically oriented branches of the sciences, such as climate and earth system research, nanostructure physics, solid-state physics, fluid mechanics, astrophysics, quantum chromodynamics, materials research, chemistry, molecular dynamics, polymer research and biophysics. Due to the incessantly progressing sophistication of mathematical models and the increasing complexity of simulations, the demand of these subject areas for computing capacity is limitless. The continuing improvement of computing power causes an increase of demand, which in turn makes new scientific questions possible and thus initiates new computational requirements. In all these areas of expertise there are applications which create a need for computing power that cannot even be met by the most powerful computing systems worldwide available today. The Wissenschaftsrat recommends the establishment of supercomputers of the highest performance class in Europe in order to provide an adequate supply of computing power to these subjects. Computers of this kind cannot be funded by one country alone; they have to be financed on a European level.

1) Wissenschaftsrat: Empfehlung zur künftigen Nutzung von Höchstleistungsrechnern. In: Empfehlungen und Stellungnahmen 2000, Cologne 2001, vol. I, p. 229-262.

Preliminary Remark

The following recommendation was compiled by the Wissenschaftsrat's national committee for the procurement and use of high performance computers. The committee is responsible for strategic advice to the Federal Government and the states in matters concerning the provision of supercomputers for science and research. To this end, the committee drafts recommendations on investment planning and carries out hearings on central issues in connection with the use and operation of supercomputers.

Some experts who are not members of the Wissenschaftsrat have taken part in the preparation of this recommendation. The Wissenschaftsrat is especially obliged to them.

This recommendation was issued by the Wissenschaftsrat on 12 November 2004 in Hamburg.

A. Computational needs of selected research areas

The availability of supercomputers for science and research is a decisive location factor in the international competition. Supercomputers are especially used in science and industry to support research where real experiments are not possible, too time-consuming or too expensive and where computing systems of lower performance levels are not adequate. They constitute an indispensable tool for the top research. In the past few years scientific breakthroughs have been achieved at numerous universities and research establishments by the use of supercomputers which would not have been possible on computers of lower performance classes. The appropriation of an adequate infrastructure is a task of research funding with public means. Open access independent of location and institutional affiliation of users is necessary.

For future applications in single research areas, the computing capacity technically available at present is however not sufficient to solve the problems at stake. The need of single research areas for computational capacity is unlimited because of the continuously progressing sophistication of mathematical models and the increasing complexity of simulations. Furthermore, the number of users and of research areas is constantly growing.

The requirements for scientific computing can be divided into the categories of capacity and capability computing. Capacity computing means a high turnover of a large number of program runs of small or medium size. This kind of computing can be dealt with comparatively cost-efficiently by providing adequate computing capacity on the regional or national level, or by the future use of grid computing methods.

Capability computing encompasses the execution of simulations using large programs. This kind of simulation imposes much greater demands on the capability of a computing system than capacity computing. Problems in the area of capability computing require a very high computing power. Calculations on a supercomputer are called capability computing if they use a significant share of processors of a large parallel computer in concert. These calculations need large amounts of memory. The

high communication bandwidth and low latency times necessary for capability computing can technically only be achieved with a single parallel computer, not with a distributed computing infrastructure. The establishment of Europe-wide supercomputer capacity is required for this class of problems.

In the following, examples of present and future research projects are outlined for which the installation of a supercomputer on a European level is mandatory.

Climate and earth system research

In the area of climate and earth system research, the demand for high performance compute services is for several reasons rapidly increasing beyond today's capacities. There is now a growing recognition that "global change" is much broader an issue than "climate change". The scientific challenges for the near future call for the development of complex Earth System Models (ESMs), extending the concept of earlier, purely physical climate models to multi-component models. These would have to account, e.g., for biogeochemical processes in the oceans, in the atmosphere, and on land. In addition, there is a growing interest in models which consider the anthropogenic component of the earth system. Models of these kinds will thus become structurally much more complex than the models in use today – with drastic consequences for their computational demands.

Next generation Earth System Models will also require much higher spatial and temporal resolutions than the present time ones. Even today, spatial model resolutions of less than 10 km and temporal resolutions of less than 20 min for the atmosphere component are being aimed at. A bisection of the resolved scales involves an increase of computation time by a factor of eight. A higher resolution promises a reduction of modelling errors because it allows the simulation of small-scale processes on the basis of fundamental physical principles. These have to be approximated in present models by more or less rough parametrisations. At the same time, a higher resolution yields more detailed spatio-temporal information and thus provides an improved decision support, e.g. in the context of land use management, or for practical applications in the insurance and re-insurance industries.

Such high-resolving models need computing systems which deliver not only top performance in specialised tests but also a high sustained performance. Their processors have to be connected by a technology allowing exchange of data with a high bandwidth and low latency time. The necessary latency time cannot be reached in a wide area network of computers. Therefore it is not possible to distribute such a model onto several spatially separated computers without a significant loss of performance.

The validation of earth system models has to be supported by the whole spectrum of available observation data since systematic experiments on the Earth System with the objective of obtaining measurement results are not possible. In the context of global climate change computational modellers intensely cooperate with palaeoclimate researchers to develop a mechanistic understanding of the climate system's workings and of earlier massive changes in the climate system. Due to the vast time spans that need to be simulated, such studies are presently possible only with relatively coarse-grained, inexpensive, but also strongly simplified models of intermediate complexity (EMICs). These do provide important insights, yet cannot yield details of many of the spatio-temporal scales of interest. Even at moderate resolution the related long-term simulations would not be feasible with today's computing resources. In order to match the models of higher resolution with such long-term data, a considerable enhancement of the time horizon of simulations is necessary which is directly proportional to the necessary computational effort.

There is a strong demand from political decision-makers to further reduce modelling uncertainties. One valid strategy to meet these demands involves running a large number of simulations that differ from each other by systematic variation of uncertain model parameters. Another similar strategy involves the comparison of model results obtained with competing computational codes when applied to the same real-life situation. Modern statistical methods then allow one to quantify the uncertainties associated with the models used and provide techniques to improve them. The parallel operation of one or several models in such ensembles entails computing time requirements proportional to the number of ensemble members. Ensembles of ten re-

alisations are minimal, ensembles of 100 to 1,000 members are needed to generate reliable statistics.

A computing power of 15 teraflops or more of sustained performance will be necessary for real-time implementation of model simulations in the future. A computer for climate models with the specified sustained performance needs a peak performance of 50 up to 150 teraflops.

Nanostructure physics

Scientific computing on supercomputers has established itself as an important pillar of solid state, surface and nanostructure physics in the last 15 years. The achieved computing power facilitated the application and the breakthrough of quantum mechanical and electron-structure based theories, which allows a realistic description and prediction of microscopic properties of real material systems without empirical input. The most important of these theories is the density-functional theory, which is solved by ab initio methods. For this purpose the many-body problem of interactive electrons in solids is represented by a one-electron problem in a self-consistent potential. This means the potential is incrementally computed. Dependent on the point and translation symmetry of the solid or the surface up to 1,000 eigenvalue problems have to be solved in each step. The computing time for each problem scales with the number of atoms in the unit cell of the solid. The computing time increases linearly to cubically with the number of atoms. The calculation has to be repeated for each external parameter, e.g. for each new adjustment of magnetisation.

Static electronic structure calculations for 200 up to 400 atoms per unit cell and ab initio molecular-dynamical calculations for simulation times of up to 10 picoseconds are possible on current supercomputers. This allows the description of complex surfaces and their reconstructions, the determination of reaction barriers for the diffusion or the analysis of chemical reactions of adsorbed molecules on metallic, semiconductive and oxidic surfaces.

Based on density functional theory, a theory for the calculation of the self-energy of quasi particles in real solids was developed. This theory allows the calculation of the

life time of electrons or the calculation of band gaps in semiconductors. In comparison with the density functional theory, the requirements for disk space and computation time are increased by a factor of 100 since the self-energy has to be calculated as a function of frequency as well as being non-local. Therefore only volume systems with up to ten atoms per unit cell have been examined up to now. This theory will only be fruitfully applied within the next few years when more powerful supercomputers become available.

Compared to nonmagnetic systems, calculations of magnetic systems may be more complex by a factor of 200 up to 500. Not until 2001 was it possible to calculate a non-collinear magnetic structure on a surface, i.e. a structure, in which the direction of the magnetisation axis changes from atom to atom. The time needed for the calculation of a complex magnetic system with about 50 atoms per unit cell is in the order of magnitude of 100,000 CPU hours. Therefore, no ab initio molecular dynamical calculations for magnetic systems presently exist which take their magnetisation dynamics into consideration although there are unique experimental results and the research field is of outstanding importance for the development of magnetic storage media.

Quantitative electron theory makes the direct theoretical analysis of experiments in solid state physics and surface physics possible. A successful example from surface physics is provided by the simulation of atomically resolved, spin-polarised Scanning Tunnelling Microscopy (STM) experiments. Experimental results of the atomically resolved Atomic Force Microscopy (AFM) were simulated less successfully since the simulation and analysis of an AFM image of a simple semiconductor surface require a total computing time of about 100,000 CPU hours. Problems of this and many other kinds can only be tackled with capability computers of a significantly higher performance.

An important trend in nanostructure physics is the shrinking of structures of already well-known materials such as the semiconductor silicon, but also of completely new materials. The importance of quantum-mechanical methods grows for the description of nano-scale structures. However, the systems are large and complex from the point

of view of quantum mechanics. The complexity of these systems is twofold: On the one hand the systems are structurally complex with many different atoms per unit cell, on the other hand the importance of the interaction between the electrons (electron correlation) grows. Therefore an important aspect of the coming years is the treatment of strongly correlated electrons with methods which exceed the density-functional. These methods combine for example the “dynamical mean field theory“ and others, which were developed in many-particle physics, with density-functional theory and allow a quantitative and material-specific description of strongly correlated electron systems.

Many further questions in the field of solid state, surface and nanostructure physics cannot or only incompletely be tackled due to a lack of computing capacity. Among these are systems of increasing complexity, such as carbon nano tubes filled with magnetic materials or molecular magnets, for which an important role is expected in future quantum electronics and storage technology.

Solid state physics

The theory of material-oriented solid state physics becomes generally more and more important. Modern solid state materials have a multiplicity of novel qualities so that manifold uses are possible even today or can be expected for the future. Possible applications of e.g., high-temperature superconductors reach from energy storage and ultra-fast electronic chips to computer tomography in medicine. Superconductors conduct electrical current without losses below their transition temperature. In this physical state of the material, electrons are bound in Cooper pairs.

High-temperature superconductors have high superconducting transition temperatures above the temperature of liquid nitrogen, which is routinely used as a coolant in engineering or medicine. Theoretical solid-state and materials physics try to understand known materials with microscopic models and to develop tailor-made new materials. They design new ways for the synthesis of functional materials in interaction with experimental physics and show new possibilities for their application. With regard to high-temperature superconductors, a basic understanding for the high transition temperatures has to be developed and the prevailing empirical search for even

higher superconducting transition temperatures and further improved material properties has to be supplemented by precise predictions.

Microscopic models are constructed taking into account experimental results, followed by an implementation on supercomputers. The interaction of the unimaginably high number of electrons (order of magnitude 10^{23} per cubic centimetre) represents a major obstacle. The high transition temperatures suggest that lattice vibrations do not bind the electrons to Cooper pairs, as generally accepted for conventional superconductors. In fact, during the past few years and with the assistance of supercomputers it could be shown that the mechanism of pairing is of a purely electronic nature, and based on the spin of the electrons. The unusually high transition temperatures can in principle be understood in this way.

A second decisive challenge has not been solved yet. Up to now it is not clear how the macroscopic superconducting current of e.g. 10^{23} Cooper pairs in a high-temperature superconducting wire comes about. The solution of the problem requires simulations for very large model systems. Finite-size scaling procedures, which are based on cluster simulations – especially on quantum Monte Carlo simulations – have been developed for this purpose. With this the limits of the currently installed supercomputers are reached since the so-called fermion sign problem of quantum Monte Carlo simulations grows exponentially with the size of the simulated model. This problem generally arises during the simulation of the behaviour of electrons (fermions) because the fermion wave function has to be anti-symmetrical concerning the exchange of two particles. However, several relatively new methods (quantum Monte Carlo simulations in connection with renormalisation-group techniques and with dynamical mean-field theory procedures) are currently tested in order to come to an microscopic understanding of pairing mechanisms and superconductivity. This comprehension is essential for replacing the still basically empirical search for improved material properties by a basic understanding which allows predictions. These calculations cannot be carried out with the currently available supercomputer generation. The successful continuation of these works in Europe depends on a further extension of computing capacity.

Fluid mechanics

The numerical simulation of fluid dynamical processes is becoming increasingly important in research. Applications can be found in almost all fields of science and technology, such as car engineering, plane and turbine engineering, environmental engineering and in weather and climate prediction. The involved fluid dynamical processes are as a rule very complex. For one, they are almost always turbulent, i.e. very small spatial and very short temporal structures occur. Furthermore, the complexity is often further increased by chemical reactions (combustions), phase transitions, heat transport or structural movements (i.e. if a solid body moves in a fluid and influences the flow, such as a vibrant plane wing), which take place simultaneously.

Accordingly, a high demand for computing capacity arises for the numerical simulation. This can be explained by means of the Reynolds number (Re), the most important parameter of fluid mechanics. The characteristic velocity of the problem enters proportionally and the viscosity of the fluids involved enters inversely proportional in this indicator. The computational effort rises strongly, approximately cubically, with the Reynolds number. Today, e.g., the flow in a miniature pump (with Re = 100) can be worked out within seconds on a standard PC while the computation time for the flow around an airplane (with Re = 10^8) would take more than 100 billion years on the currently most powerful supercomputers presently available anywhere. From this it can be deduced that the requirements for computing capacity for flow simulations are unlimited.

In order to carry out calculations for practically relevant problems at all, which have almost always have a very high Reynolds number, it is necessary to reduce the complexity of the problem in advance by simplifying the model assumptions. Such simplifications (e.g. turbulence models, chemistry models) are however only possible to a certain degree with regard to the quality of the calculated results, so that the reduced problems also still have a high complexity. Detailed simulations with an enormous need for supercomputer capacity are indispensable for the development and validation of simplified models.

The availability of adequate supercomputer capacity for many current fluid mechanical issues is a prerequisite for conducting internationally competitive research and development. This also has to be considered against the background of an intense development of the methods of scientific high performance computing during the past few years, providing powerful and highly developed algorithms for the fluid dynamical calculations on supercomputers even for nowadays use.

Astrophysics

In astrophysics and gravitational physics, a great demand for supercomputer capacity is expected for the years ahead. Laser interferometers as gravitational wave detectors have recently been built at great expense in several locations worldwide. Despite their astounding sensitivity of 10^{-21} for relative length changes, most signals can only be extracted from the noise if their shape is known beforehand. The strongest type of source, likely the first to be detected, is the merger of a binary system consisting of two neutron stars or black holes. Calculating the precise shape of the radiated gravitational waves requires the three-dimensional and time-dependent numerical solution of the Einstein field equations - mathematically and numerically an extremely demanding problem which is tackled by research groups all over the world.

The numerical calculation of the time development of the spatial structure of the cosmos also requires enormous calculating power and storage capacity. This problem requires the numerical simulation of the internal dynamics of billions of galaxies and their interactions, including the influence of dark matter. Simulating processes of accretion onto compact objects such as white dwarfs, neutron stars and black holes is another important question in numerical astrophysics. These are three-dimensional, time-dependent problems involving hydrodynamics and magnetohydrodynamics. If turbulence is to be included, they cannot be solved on any supercomputer currently installed.

A further challenge for numerical astrophysics is the formation of stars and planets. The gravitational collapse of molecular clouds has to be modelled on very different length scales. The simulation of non-linear and turbulent processes happening simul-

taneously can only succeed if computer programs optimized for this task and next-generation supercomputers become available.

Supernova explosions and gamma bursts mark the end of the development of a star. They are the events in the universe involving the highest amount of energy; the heavy chemical elements are formed in these processes. Complex nuclear and particle physical reactions as well as transport processes determine the hydrodynamic and magnetohydrodynamic evolution which leads to the explosion of a star. Computer simulation of these processes, involving multi-dimensional cosmic ray hydrodynamics, are among the most demanding challenges of modern astrophysics and require the most powerful supercomputers available. However, only the use of supercomputers of the next generation will allow the prediction of observable quantities, such as explosion energy and distribution of elements, with higher precision than possible now.

Quantum chromodynamics

In the context of the standard model of elementary particle physics, quantum chromodynamics (QCD) is regarded as the fundamental quantum field theory of the strong interaction. It describes the components of the atomic nucleus – protons and neutrons – as binding states of three elementary particles, the quarks. The strong interaction links the quarks together by the exchange of gluons, similar to the exchange of photons causing the electric attraction by which electrons and nuclei combine into atoms. In analogy to electric charge, the interaction is described by strong charge.

The formulation of QCD on a four-dimensional space-time-grid as well as its numerical-statistical simulation on supercomputers is called lattice quantum chromodynamics (lattice QCD). The computer simulation is the only known procedure up to now to evaluate the QCD ab initio just in the energy range where the transition of quarks and gluons from quark-gluon-plasma phase to a proton and a neutron takes place. Due to quantum fluctuations, the strong charge acquires a high value at low energies, and as a consequence, perturbative calculations do no longer work.

Nuclear and elementary particle physics meet precisely in this energy range with the common interest in understanding the perceptible nuclear forces from the first principles of strong interaction. Ab initio means in this context that a fundamental theory such as QCD has to be evaluated without the detour of models or approximations, if one wants to confirm it or to use its full predictive power.

Lattice QCD has matured in the last ten years and has become an indispensable instrument for the solution of theories of nuclear and elementary particle physics as well as for the planning, evaluation and interpretation of existing and future accelerator experiments. Fundamental quantities of elementary particle physics, such as the charge of a strong force or the masses of quarks, can be determined only with the help of lattice QCD. Apart from the prediction of exotic matter, such as the so-called glueballs, or the determination of distribution functions of quark or gluon densities, the lattice QCD is regarded as the only method which can lead to an understanding of the properties of the quark gluon plasma.

The German research groups, together with those from United Kingdom and Italy in Europe as well as the USA and Japan, are leading in the worldwide activities in the field of lattice QCD. Whether European groups will be successful in the future and stay at the top of the world depends entirely on the availability of adequate computer resources.

Meanwhile, the solution of lattice QCD is considered one of the grand challenge problems of computational science. Lattice QCD is a striking example for the use of capability computing because the four-dimensional structure of the problem on the one hand requires an extraordinarily amount of computing power but on the other hand guarantees a good scalability. In doing so, a distinction can be made between the stochastic simulation of the theory, where ensembles of gluon-field-configurations are generated, and the evaluation of these fields for the determination of physical values.

This evaluation can lead to very complex program structures with often very large input-output rates as well as requirements for data management and computing po-

wer. However, the stochastic simulation is far more costly. Since about ten years ago, shared computing systems have not been sufficient any more to perform this task. The medium-term requirement for computing power adds up to about 15 teraflops of sustained performance.

Considering the integration of user groups from the new member states of the European Union, the idea of supercomputers on a European level would contribute decisively to a cooperation in stochastic simulation of QCD, analogous to the operation of large scale accelerator facilities at GSI (*Gesellschaft für Schwerionenforschung*), DESY (*Deutsches Elektronen-Synchroton*) and CERN (*Conseil Européen pour la Recherche Nucléaire*). The subsequent evaluations could be carried out on national mainframes and supercomputers, which are linked in a grid-network anyway. Here the project DEISA (Distributed European Infrastructure for Supercomputing Applications), financed by the European Union, should be mentioned.[2)] In this way, European physicists could take over worldwide leadership in the field of lattice QCD.

Materials science

Electronic structure theory and statistical mechanics belong to the most important pillars of physics and chemistry. The properties of many relatively simple materials can be quantitatively described and also physically understood by applying these theories. Compared with that, there is still a considerable lack of comprehension of the manifold phenomena in more complex and also technically more important materials. The reason for this is that time and length scales on which the qualities of each material emerge and on which the processes within the material operate vary over many orders of magnitude. This pertains to - among others - electronic, magnetic and optical components as well as sensors and catalysts. Further examples for research requirements in the field of materials research are surface coating and friction, e.g. applicable to corrosion prevention and scratch resistance. In all these areas there are close links to nanostructure physics.

---

2) Cp. paragraph „EU project DEISA“ in chapter C „European supercomputers“.

The field of electronic structure theory with the above-sketched material-scientific objective is in a very active phase, which is characterised by fast development of fundamental theories as well as of new methods, algorithms and computer programs. A theory for atomistic modelling of complex material systems has emerged during the past few years, which allows concrete predictions on the basis of the quantum theory (especially of the density functional theory) of the interacting multi-electron problem. These ab initio theories include the reliable description of the involved elementary processes, such as the breaking and building of chemical bonds, as well as the statistical mechanics of the interaction of the processes involved. Current and sustainable fields of application for this are studies on the activity of catalysts and on crystal growth. Further prominent research themes comprise the field of opto-electronics, electronic transport and friction. Decisive for an adequate description is the ab initio calculation of excited electronic states. Important methods, which need to be further extended, are the calculation of the self-energy (GW methods with vertex adjustment and the Bethe-Salpeter equation for the description of correlation effects, respectively), improved band structure methods ("exact change“ potential) as well as time-dependent density functional theory.

Important simulations from materials sciences on supercomputers concern themselves with the modelling of catalysis und crystal growth, e.g., the growth of nanostructures and self-organised semiconductor quantum dots on semiconductor surfaces. The computing requirements of a typical catalysis project would be approximately 50 years on a processor with about six gigaflops peak performance today. Thus the same calculation would take two years on 25 processors. This applies only for the simulation of a simple model of a catalyst material with one simple chemical reaction, e.g., for the oxidation of carbon monoxide on an oxidated ruthenium surface. Calculations on semiconductor quantum dots, on materials of spintronics, on isolators with high dielectric constants ("high-k-dielectrics“) for electronic components or for the stability of bio-molecular structures have similarly high computing requirements. More complex processes on surfaces, e.g., the conversion of hydrocarbons on variable metal-oxide surfaces or nanoparticles, cannot be calculated even today. Similarly, the calculation of dynamical processes that determine e.g., the transport properties of semiconductor structures is still not computationally feasible yet. Such

examinations require an at least ten times higher, for comprehensive studies a 100 times higher computing capacity than currently available.

The next aim is to carry out investigations on a single processor within a month where currently still 50 CPU years are needed. Only then the results will be available sufficiently fast to compare several systems and to identify trends. This would mean a big step forward for basic research as well as for the application oriented materials sciences. The most costly calculations refer to electronic structure based ab initio simulations of the mentioned problems from physics and chemistry as well as force field based simulations of bio-molecular processes.

Theoretical chemistry

The non-experimental examination of complex molecular systems, i.e. of interacting quantum mechanical multi-particle problems, is at the centre of interest of modern theoretical chemistry. Molecules and clusters, but also liquids, surfaces, chemical sorbates and bio-molecules belong among these systems. Generally, the interest in investigations on bio-molecules is rising.

An efficient procedure for calculating the electronic structure as well as the dynamics of molecular multi-particle systems on an atomic scale is the method of ab initio molecular dynamics developed by Car and Parrinello. Two extensions of this are constituted by the fully quantum mechanical ab initio path integral method and the non-adiabatic generalisation. The first extension includes important quantum mechanical effects such as tunnelling and zero-point vibrations, which can be essential especially for hydrogen, e.g., during proton diffusion. The non-adiabatic extension is able to examine couplings between several electronic states beyond the Born-Oppenheimer approximation and thus makes photo-chemistry accessible, e.g., radiation damage of DNA.

With such ab initio simulations, the computer together with efficient algorithms and programs becomes a virtual laboratory, in which chemical, physical and biochemical processes run dynamically "in silico“ on the fundamental level of electrons and atom cores. Using the standard Car-Parrinello method one has been able to study time-

resolved dynamical processes without experimental knowledge from nature for about ten years. The current development in the field of ab initio simulations is dominated by issues from the field of materials, nano and especially life sciences, which become more and more complex.

Such simulations, which are based on the fundamental laws of physics and chemistry, require a considerable use of computer resources. For reasons of efficiency, the quantum mechanical wave function of the total system has to be kept in central memory permanently. This requires primary storage on the order of 50 gigabytes if the standard Car-Parrinello method is used for today's standard problems (about 100 heavy atoms). A typical simulation run of about ten picoseconds needs several CPU months on today's computers.

The requirements on computing power and memory consumption dramatically rise with the size of the system. The world record is held by a 5 picoseconds simulation of a water surface consisting of 250 molecules, for which a computer with a peak performance of 11 teraflops was used for some months. Due to lack of computing resources, such calculations are not possible in Europe even though the software was developed in Europe. Such simulations will become standard applications (e.g., after the installation of Blue Gene/L) so that "cutting edge" problems will then lie in the region of 100 to 500 teraflops peak performance. It will then be possible to treat small enzymes with the standard Car-Parrinello method, which would open the era of quantum biology of dynamical catalytic processes.

For ab initio path integral simulations, in which the atomic cores are also treated quantum mechanically, the requirements on memory consumption and computing time rise by one to two orders of magnitude relatively to the standard Car-Parrinello method previously discussed. However, these calculations easily lend themselves to a hierarchical parallelisation due to their mathematical structure. For non-adiabatic Car-Parrinello simulations, the running time (time step: factor 10 smaller, Monte Carlo averaging over about 10 to 100 independent molecular-dynamic trajectories) increases by two or three orders of magnitude. Thus a peak performance of one

petaflop would be necessary for the parameter-free examination of nano-catalytic and biochemical processes.

Molecular dynamics

In classical molecular dynamics, the paths (trajectories) of all single particles of a multi-particle system are calculated under the influence of mutual interactions and external forces. The trajectories of all particles are derived from the numerical integration of the equations of motion so that any observable can be determined. In a sense, a numerical experiment is carried out, the progress of which in contrast to a real experiment is strictly controllable. Since the fundamental physical equations are used directly, it is not necessary to model the examined processes beforehand in a simplified way. Thus molecular dynamics is used especially where modelling is not at all possible.

Only the underlying laws of motion have to be modelled. Classical effective interactions are sufficient for many systems. With this, very large systems of several million particles can be simulated. For some purposes however, e.g., for taking into account chemical reactions, the electrons have to be described quantum mechanically with ab initio molecular dynamical methods. This considerably limits the system sizes which can be treated on today's computers.

Molecular dynamics is a universal method with a very broad application spectrum. It is used for the examination of properties of complex solids and polymers in the materials sciences, of phase diagrams in process engineering, of the structure of macromolecules for the optimisation of active ingredients in pharmaceutical chemistry, the folding and dynamics of bio-molecules, the pouring behaviour of granular media such as flour and sand, of shock waves in solid bodies, liquids and gases, the movement of galaxies in the cosmos and much more. The behaviour of such multi-particle systems is very complex and hardly approachable for analytic theory. However, molecular dynamics can especially serve for uncovering the fundamental processes in the system and to deliver the material parameters which can then be used in further-reaching models for engineering calculations.

Simulations on the level of single particles involve a considerable computational cost. They take about ten microseconds per interval and particle on modern processors. On the atomic scale, an interval of one femtosecond has to be simulated to achieve a precise scanning of atomic vibrations. If the movements of ten million particles are to be calculated for a one nanosecond interval, this means a simulation time of 10^8 CPU seconds, equalling approximately 30,000 CPU hours or 1,200 CPU days, is required. Hence, supercomputers are needed even for the comparatively short simulation time of one nanosecond. Those simulations are in fact indispensable in the materials sciences if mechanical processes such as crack growth or material fatigue are to be simulated.

Soft Matter

Soft matter research comprises physics and chemistry especially of synthetic and biological polymers, colloids, membranes as well as organic-anorganic hybrid systems. Classical chain molecules, i.e., polymers in the narrower sense, meanwhile only form a subgroup of soft matter and serve as a reference for model building.

Thermal fluctuations play an important role because of the relatively low energy density in these materials: the materials are “soft“. The spatially large molecules are able to strongly fluctuate in their form (conformation). As a consequence, processes on the microscopic-atomic, and those on the mesoscopic level respectively, contribute to the materials’ properties in equal measure. There are often ten or more orders of magnitude between the typical time scales of the local atomic movements and the meso- respectively macroscopic phenomena. More than 3,000 CPU years are needed on standard processors if the computing time for the molecular dynamical simulation of an atomically resolved polymer melt (e.g., polystyrene) is estimated on the basis of a relatively small system of 100,000 atoms and a typical relaxation time of 10^{-5} seconds. But then only the behaviour of one system during one relaxation period would be simulated. Questions concerning the behaviour under shearing strain (e.g., extrusion during the fabrication of CDs by injection moulding) or concerning varying polymerisation degrees require many simulations of this kind. Therefore simulations of the dynamics of polymer melt can only be conducted on a strongly idealised level. Similar considerations apply for the simulation of simple polyelectro-

lytes in implicit solvents (e.g., water is considered only as a dielectric homogeneous background), of multi-scale simulations of systems with relatively short chains or for studies on multi-component systems. Other scientifically interesting questions are even more complex and time-consuming. But despite all restrictions, there has been impressive progress in the last few years.

The desirable permanent monitoring of experimental studies will only be possible if the computer performance available today is increased dramatically. Further parallelisation cannot be the solution since the systems' temporal development has to be followed in many cases. Therefore there are, apart from the need for considerably more powerful supercomputers, intensive efforts to develop simulation methods with which several length and time scales can be systematically linked to each other (multi-scale simulations). A real progress is only possible if both developments go hand in hand. Examples are the consideration of local ion interactions and explicit solvents (e.g., molecular structure of water) for polyelectrolytes, to which almost all biopolymers belong, the dynamics of realistic polymer melts with branched polymers, the phase behaviour of multi-component systems or scale-spanning calculations with realistic dynamics and with conformation changes of smaller and later also of larger biopolymers.

Supercomputers that are by orders of magnitude faster than the ones available today would allow access to important new areas of investigation. These include e.g., suppression of turbulence in liquids by addition of polymers ("turbulent drag reduction“), membrane fluctuations and membrane functions including their interaction with membrane proteins and the realistic consideration of the surrounding water, the structure building and function of molecular aggregates as well as the connection of the variable conformation of macro-molecules with functional groups (e.g. chromophores in fluorescent polymers) and their electronic properties. The last examples are closely related to other fields of computational science, e.g. catalysis research, quantum chemistry and fluid mechanics. In order to stay competitive in the aforementioned fields and to take part in the described developments, the available computational power has to be by orders of magnitude larger than today and the access to it on national and international level is indispensable.

Biophysics (Life Sciences)

The computer simulation of the dynamics of bio-molecules describes their atomic movement by numerical solution of Newton's classical laws of motion for all atoms involved. In this way a movie of the functional actions, e.g. in a membrane protein, is generated. A typical simulation system currently comprises a protein which is embedded in a lipid membrane. Such a complex consists of about 100,000 atoms whose mutual $5x10^9$ inter-atomic interactions have to be taken into account completely because of their long range character. For an integration time step of typically 10^{-15} seconds, $2x10^{17}$ floating point operations are needed in order to describe the dynamics of the system over the relatively short time span of one nanosecond.

In light of the fact that most biochemical processes take place on much longer time scales than a nanosecond – in the case of protein folding approximately between microseconds and seconds – most biological questions still can not be answered because of limited available computing power. Especially reliable simulations of large systems such as the ribosome (about $2x10^{16}$ atoms) and of processes which require very long simulation times such as protein folding can only be performed with the help of more powerful supercomputers. Each increase in available computing capacity will facilitate new biological insights and perceptions on molecular-biological mechanisms and processes with potentially pharmacological impact. Accordingly, parallel supercomputing is a central instrument of the computer-supported and theoretical biophysics, and hence is indispensable for this field.

In addition to these requirements of conventional molecular dynamics, additional demand for computing time arises owing to the increasing use of methods of statistical mechanics ("ensemble molecular dynamics"), which require the parallel calculation of a large number of simulation trajectories, as well as by the increasing demand for first-principles-simulations of experiments. To fulfil this demand, a wide spectrum of supercomputers with different balancing and optimisation between communication and computing power is mandatory, as e.g., defined by the "Thunder"-installation in the Lawrence Livermore National Laboratory on the one hand and the Blue Gene/L or the Earth Simulator on the other hand. Capability computing, therefore, has a central impact for biophysics and life sciences in Europe.

Starting from the sequence (“human genome project“), and continuing with the dynamics and function of bio-molecules, the life sciences will focus more strongly on their interactions and networking (systems biology) with the aim to quantitatively understand complete metabolic function networks and thus e.g., to be able to dose drugs better and more individually. Considering the large number of proteins involved (about 28,000 for the human organism), folding simulations on peptides or very small proteins, which presently require considerable efforts (several months of wall clock time on a supercomputer), will have to be coped with as a matter of routine in large numbers in the near future. Of special importance is the understanding of enzyme catalysis, which requires the combination of force-field based and quantum chemical codes and further increases the computing power requirements by two to three orders of magnitude.

B. Available and planned computing power

National computing power

Among the 500 most powerful mainframes and supercomputers worldwide,[3)] 35 computers are installed in Germany. Nine of them are available for universities and extra-university research institutes, the other 26 are commercially operated and used by companies. The ranking of the installed computers in the TOP500 list is mainly oriented toward the theoretical peak performance, so that conclusions about the really attainable performance (sustained performance) for scientific applications can only be drawn within limits.

The currently most powerful supercomputer in Germany is installed in the John von Neumann Institute for Computing at the Research Centre Jülich. It has a peak performance of 8.9 teraflops and was put into operation in February 2004. It ranks 30th in the TOP 500 list from November 2004.

In spring 2005 a computer with a peak performance of 11 teraflops will be installed in the High Performance Computing Center in Stuttgart.

The state of Bavaria is going to procure a computer with a peak performance of 20 teraflops in the new building of the Leibniz Computer Centre in Garching in 2005,[4)] which will be expanded to 40 teraflops in 2007. Thus the development of supercomputing in Germany follows the Wissenschaftsrat's recommendation to provide at least one system of the highest performance class on the national level every two to three years.

Apart from the pure hardware, the support for users in the usage of supercomputers is of special importance. This support has evolved on a broad basis in Germany. The hierarchic provision concept with the three national top-level centres: High Performance Computing Center in Stuttgart, John von Neumann Institute for Computing

---

3) TOP 500 Supercomputer Sites: www.top500.org, November 2004.

4) Wissenschaftsrat: Stellungnahme zur Anmeldung des Landes Bayern auf Neubau eines Gebäudes für das Leibniz-Rechenzentrum und zur Beschaffung eines Höchstleistungsrechners im Zusammenhang mit der Errichtung eines Gebäudes für das Leibniz-Rechenzentrum zum 32. Rahmenplan. In: Empfehlungen und Stellungnahmen 2002, Cologne 2003, vol. II, p. 83-102.

Jülich und Leibniz Computer Centre in Garching, as well as several high performance computing centres, accompanied by the funding of networks of competence, is unrivalled in Europe. This development status is a result of intense funding of supercomputers and high-performance systems in Germany. The computing science in the centres and in universities provided the basis for an adequate use of supercomputers and high-performance systems via an interdisciplinary cooperation with the specialised sciences targeting the applications, by the development of specific software for these types of computer and by the exploration of fundamental properties of various computing architectures.

The German infrastructure for application support is also competitive in comparison to American research centres such as the ones of the National Science Foundation. However, if one takes into account the centres operated by the Department of Defence and the Department of Energy respectively,[5)] these still have a head start because of their size, especially in the development of prototypes for data management and in the development of algorithms as well as in the conception of new paradigms (e.g., grid computing).

Computing power in Europe

128 of the worldwide 500 most powerful high performance systems and supercomputers are located in twelve European countries – in Belarus (1), Germany (35), Finland (1), France (15), Italy (15), the Netherlands (5), Portugal (3), Russia (1), Sweden (4), Switzerland (3), Spain (3) and the United Kingdom (42). At least 35 of these computers are used in the context of science and research.

In the European comparison of computing power of high performance systems and supercomputers, Germany has for many years shared ranks in the top positions with the United Kingdom as far as the peak performance of the systems as well as the accumulated performance of the systems is concerned. The fastest computer in Great Britain holds rank 11 of the most powerful supercomputers worldwide. One other British computer ranks among the 20 fastest computers. Germany is ranked

5) E.g. Los Alamos National Laboratory, Lawrence Livermore National Laboratory, National Energy Research Scientific Computing Center.

30th among the TOP 500 with the installation at the Research Centre Jülich in February 2004.

The supercomputer HECToR (High End Computing Terascale Resources) with a peak performance of 100 teraflops will be installed in Great Britain in 2007. Further goals are a doubling of computing performance to 200 teraflops until 2008, and after that to 400 teraflops until 2010.

The currently most powerful European supercomputer was installed in the new Barcelona Supercomputing Center at the Technical University of Catalonia in October 2004. It has a peak performance of 31.3 teraflops and holds rank 4 in the TOP 500 list. The costs are estimated at 70 million Euro.

International computing power

In a worldwide comparison it can be stated that significantly more powerful systems were installed in the USA and Japan (Appendix 1). The support of the local computer industry played a decisive role in both countries. In the US, the use of the largest systems was reserved for the Department of Defence and the Department of Energy respectively for many years. These systems have for the most part been open also for civil projects since about four to five years ago. Furthermore, the three supercomputing centres of the National Science Foundation[6)] possess supercomputers which are used solely for civil research and science.

The most powerful supercomputer worldwide is located in the US. It has a peak performance of 91.8 teraflops. Eight of the ten fastest supercomputers worldwide are located in the US, one in Japan and one in Spain. For these supercomputers, performances (LINPACK values) of 9.8 to 70.7 teraflops and theoretical peak performances of 15.3 to 91.8 teraflops are specified (Appendix 1).

---

6) These are the Center for Supercomputing Research & Development, Illinois (www.csrd.uiuc.edu), the Pittsburgh Supercomputing Center (www.psc.edu) and the San Diego Supercomputer Center (www.sdsc.edu).

In the US, a plan for high-end computing was presented by the Office of Science and Technology Policy of the White House in May 2004.[7)] It recommends a closer cooperation concerning the use and development of resources between the federal agencies which have access to supercomputers. Furthermore, it points out that some institutions have difficulties in obtaining the computing capacity necessary for large research projects, especially if the calculations can only be carried out on supercomputers. The Department of Energy intends to build a supercomputer at the cost of 250 million US Dollar with a peak performance of 250 teraflops. This computer shall be erected at the Oak Ridge National Laboratory in Tennessee.

At the beginning of the year 2005, the “Blue Gene“ computer with a peak performance of 360 teraflops shall be installed in the USA, priced at 100 million US Dollars. This supercomputer is being built for the Lawrence Livermore National Laboratory in California. Follow–on machines are planned for the years 2006/2007 (“Blue Gene/P“) and 2007/2008 (“Blue Gene/Q“) and shall have a peak performance of up to 1,000 respectively 3,000 teraflops.

7) www.hpcc.gov/pubs

C. European supercomputers

A continuous qualitative and quantitative upgrading of the computing resources is indispensable for the safeguarding of scientific competitiveness. The acquisition of supercomputers of the highest performance class is necessary to keep Europe and thus Germany in the competition with Japan, the US and China. The costs for internationally leading supercomputers add up to 200 million euro and can no longer be paid by a single country. The competitive ability in the international context could be reached much more efficiently in the long run by European funding and the pooling of national resources, as is already the case for scientific and technical large-scale facilities. For this reason, the Wissenschaftsrat recommends the establishment of supercomputers of the highest performance class on a European level.

EU-Project DEISA

The acquisition and usage of European supercomputers requires some co-ordination between the countries involved in the establishment of such computing centres. The project DEISA[8)] (Distributed European Infrastructure for Supercomputing Applications), which is funded by the European Union in the context of the 6th Frame Programme, can be seen as an essential preparation for the establishment of European supercomputers. This project was started on the 1st of May 2004 with a budget of 24 million euro for the duration of five years. High performance and supercomputing centres from Germany, Finland, France, Great Britain, Italy and the Netherlands[9)] are participating in DEISA, Spain will fully join in the future. It is the intention of the project to build up a network of existing supercomputers in the European states and to operate them jointly. To this end, the existing high performance systems and supercomputers of the European countries involved in the project are linked together. An important prerequisite for an efficient linking is the transparent access to data, which shall be achieved by a global data system for all supercomputing centres. The thus

---

8) www.deisa.org

9) Present members: Germany: Research Centre Jülich, Computer Centre Garching of the Max Planck Society; Finland: Finnish Information Technology Centre for Science; France: *Institut du Développement et des Ressources en Informatique Scientifique*; United Kingdom: Edinburgh Parallel Computing Centre, European Centre for Medium-Range Weather Forecasts; Italy: *Consorzio Interuniversitario*; Netherlands: SARA Computing and Networking Services. Future members: Leibniz Computer Centre in Munich, Computer Centre Stuttgart, Spanish High Performance Centre in Barcelona.

combined computing power will add up to more than 100 teraflops at the beginning of the project. The project does not provide any additional computing power; instead it rather serves to make computing resources efficiently available for outstanding European scientists and to open up the possibility to use the strongest European computer in each case. For this reason the DEISA-project can be regarded as the ideal starting point for the establishment of European supercomputers.

ERA-Net

The Wissenschaftsrat welcomes the beginning understanding of science institutions in Europe for the provision of computing capacities on a European level since the competitive ability in the international context can be permanently achieved only by a European initiative and the pooling of national resources. Germany, the United Kingdom and France have joined forces and are currently preparing an application for an ERA-Net (European Research Area) for funding by the European Union. The ERA-Net shall serve to make preparations for the establishment and the operation of one or several European supercomputers. The establishment of collectively used European supercomputer centres is the only possibility to meet the need for highest computing power in Europe. If the European countries which are qualified for top level systems pool their resources, personnel and infrastructure, very much larger and more powerful systems can be installed than would be possible on the national level, considering the present and medium-term budget situations.

European computer pyramid and acquisition strategy

The present provision for computing capacity in Germany consists of four levels: 1. the work station level, 2. the institute/department level, 3. the university level of the high performance computing centres as well as 4. the level of supercomputers available in this country. In the Wissenschaftsrat's opinion, the computer pyramid in Germany has proved itself and should be expanded by an additional level of European supercomputers. This means that supercomputers which are more powerful than the national supercomputers should be installed in Europe. The supercomputers available within Europe should be cyclically replaced in regular intervals by new, more powerful systems. The purchases should be coordinated so that in the interval of relevant innovation cycles – approximately every two to three years – a state-of-the-

art system of the highest performance class is available for European science. This means in the long run that the establishment of three European supercomputers is necessary with an operating lifetime of a computer of about five to six years.

It is a condition for the extension of the computer pyramid on a European level that the national high performance and supercomputing centres retain their function in the computer pyramid and are equipped with adequate computers on the previous level in regular intervals. European supercomputers represent a complement, not a replacement of the national supercomputer centres.

Location

Each computing centre in Europe presently operating on the national level should be allowed to apply as the location of a European supercomputing centre. The establishment of new centres in the open countryside is not expedient in the Wissenschaftsrat's opinion since, according to experience, it takes the time of a whole generation of scientists to achieve the complete operability and professional competence for a centre of high quality. In addition, the Wissenschaftsrat argues for the connection of a European computing centre to a university or an extra-university research establishment according to its recommendation for the national provision with computing capacity. In the Wissenschaftsrat's opinion it stands to reason that Germany should be the location of one of the European supercomputing centres because of its long standing experience in supercomputing, its established user support and its already installed centres of competence.

Allocation of computing time

Computer time should be allocated under scientific considerations, not proportionally according to countries. In this context the Wissenschaftsrat recommends to orient oneself towards European centres such as CERN (*Conseil Européen pour la Recherche Nucléaire*) in Switzerland or the ECMWF (European Centre for Medium-Range Weather Forecasts) in the United Kingdom. In these centres, the access to the resources is allocated by scientific steering committees. The computing capacity of a European supercomputer should only be granted for tasks which cannot be car-

ried out in the national high performance and supercomputing centres or only with unjustifiable expenses.

Networking

An effective networking of European supercomputing centres with the national high performance and supercomputing centres of the European countries has to be ensured. Prerequisites are the provision of a sufficient network transmission capacity, appropriate data transmissibility as well as ensuring the necessary data security. This should be supplemented by a harmonisation of the organisational and administrative conditions at the computer locations.

Service performance

Apart from planning, installation and operation of the central machine resources, a European supercomputing centre has to render extensive services for top level research. The Wissenschaftsrat regards it as necessary to resort to a network of professional competence of computer-related advisory services of methodical-professional and mathematical-informatical kind.

Software development

Software development which aims at an increased performance of supercomputers is of special importance. Progresses in the solution of big problems are never made by the installation of high-performance hardware alone but always also by progresses in software development. European computing centres have to function as crystallisation points for the formation of networks of competence and computational science communities. Apart from the operation of the supercomputer itself it has to be checked which further tasks European supercomputing centres are supposed to carry out and which general functions they shall have, especially in the development of software and algorithms, in know-how transfer and in the area of education as well as extended training.

Funding

The funding of a supercomputer of the highest performance class in Europe should be provided in concert by the countries involved in the acquisition of the system. The

cost for a supercomputer currently amounts to about 200 million euro.[10)] A joint funding in Europe stands to reason if one considers the cost related to the coming increase in computing power, as can be observed in the US.

Germany has built up a very good infrastructure in the last years, considered to be among the best in Europe. For this reason, the Wissenschaftsrat considers Germany as an excellent location for such a computer. If the assignment of a European supercomputer to Germany can successfully be performed, the Federal Government and the states would have to pay costs in the region of 30 to 50 million euros. This is approximately the amount spent for a national supercomputer.

The Wissenschaftsrat asks the Federal Government to work towards the establishment of a supercomputer in the committees of the European Union and to ensure that Germany will be the location of a European supercomputer.

[10)] The Japanese supercomputer “Earth Simulator“, a pilot project, cost 400 million US Dollars. In the case of a European supercomputer, it can be assumed that the costs are significantly lower.

D. Appendix

Appendix 1: List of Top10 Supercomputers from the TOP 500-List (November 2004)

Rank	Computer Manufacturer	Linpack-value in gflops	Peak per-formance in gflops	Location Country/Year	Field of Application
1.	BlueGene/L DD2 beta-System (0.7 GHz PowerPC 440) / 32768 IBM	**70720**	**91750**	IBM/DOE United States/2004	Research
2.	SGI Altix 1.5 GHz, Voltaire Infiniband / 10160 SGI	**51870**	**60960**	NASA/Ames Research Center/NAS United States/2004	Research
3.	Earth-Simulator / 5120 NEC	**35860**	**40960**	The Earth Simulator Center Japan/2002	Research
4.	eServer BladeCenter JS20 (PowerPC970 2.2 GHz), Myrinet / 3564 IBM	**20530**	**31363**	Barcelona Supercomputer Center Spain/2004	Academic
5.	Intel Itanium2 Tiger4 1.4GHz - Quadrics / 4096 California Digital Corporation	**19940**	**22938**	Lawrence Livermore National Laboratory United States/2004	Research
6.	ASCI Q - AlphaServer SC45, 1.25 GHz / 8192 HP	**13880**	**20480**	Los Alamos National Laboratory United States/2002	Research
7.	1100 Dual 2.3 GHz Apple XServe/Mellanox Infiniband 4X/Cisco GigE / 2200 Self-made	**12250**	**20240**	Virginia Tech United States/2004	Academic
8.	BlueGene/L DD1 Prototype (0.5GHz PowerPC 440 w/Custom) / 8192 IBM/ LLNL	**11680**	**16384**	IBM - Rochester United States/2004	Vendor
9.	eServer pSeries 655 (1.7 GHz Power4+) / 2944 IBM	**10310**	**20019.2**	Naval Oceanographic Office (NAVOCEANO) United States/2004	Research
10.	PowerEdge 1750, P4 Xeon 3.06 GHz, Myrinet / 2500 Dell	**9819**	**15300**	NCSA United States/2003	Academic

Appendix 2: List and position of all supercomputers installed in Germany from the TOP 500-List (November 2004)

Rank	Computer Manufacturer	Linpack-value in gflops	Peak per-formance in gflops	Location/Year	Field of Application
30.	eServer pSeries 690 (1.7 GHz Power4+) / 1312 IBM	**5568**	**8921**	Research Centre Jülich/2004	Research
50.	BladeCenter Xeon 3.06 GHz, Gig-Ethernet / 1064 IBM	**3755**	**6511.68**	Bank (H)/2004	Industry *Finance*
85.	pSeries 690 Turbo 1.3 GHz / 822 IBM	**2198.44**	**4274.4**	Max Planck Society - IPP/2003	Research
93.	SP Power3 375 MHz 16 way / 1920 IBM	**2106**	**2880**	Deutscher Wetterdienst/2003	Research *Weather and Climate Research*
94.	ALiCEnext, Opteron 1.8 GHz, GigE, Parastation / 1024 Angstrom	**2083**	**3686.4**	Wuppertal University/2004	Academic
154.	Opteron 2.0 GHz, GigE / 896 NEC	**1778**	**3584**	DaimlerChrysler/ 2004	Industry *Automotive*
168.	SR8000-F1/168 / 168 Hitachi	**1653**	**2016**	Leibniz Computing Centre/2002	Academic
181.	DL140, Opteron 2.2 GHz, Myrinet / 512 HP	**1576**	**2252.8**	AMD/2004	Industry *Semiconductor*
182.	DL140, Opteron 2.2 GHz, Myrinet / 512 HP	**1576**	**2252.8**	AMD/2004	Industry *Semiconductor*
198.	SX-6/192M24 / 192 NEC	**1484**	**1536**	DKRZ - Deutsches Klimarechen-zentrum/2003	Research *Weather and Climate Research*
214.	pSeries 690 Turbo 1.3GHz / 512 IBM	**1384**	**2662.4**	HLRN at Hannover University - RRZN/2004	Academic
215.	pSeries 690 Turbo 1.3GHz / 512 IBM	**1.384**	**2662.4**	HLRN at Zuse Institute Berlin/2004	Academic
236.	SuperDome 875 MHz/HyperPlex / 640 HP	**1263.3**	**2240**	Deutsche Telekom AG/2004	Industry *Telecomm.*
249.	xSeries Xeon 3.06 GHz, Myrinet / 260 IBM	**1231.05**	**1591.2**	DaimlerChrysler/ 2004	Industry *Automotive*
256.	xSeries Xeon 3.06 GHz, Myrinet / 256 IBM	**1215.31**	**1566.72**	MTU Aero Engines/ 2004	Industry *Aerospace*
260.	Integrity Superdome, 1.5 GHz, HPlex / 320 HP	**1210**	**1920**	HP Financial Services/2004	Industry
266.	Integrity Superdome, 1.5 GHz, HPlex / 288 HP	**1210**	**1728**	SDH/2004	Industry

Rank	Computer Manufacturer	Linpack-value in gflops	Peak per-formance in gflops	Location/Year	Field of Application
273.	Integrity Superdome, 1.5 GHz, HPlex / 256 HP	**1210**	**1536**	SARL International/ 2004	Industry
287.	Integrity Superdome, 1.5 GHz, HPlex / 256 HP	**1198.1**	**1536**	SARL International/ 2004	Industry
288.	Integrity Superdome, 1.5 GHz, HPlex / 256 HP	**1198.1**	**1536**	SHD/2004	Industry
347.	BladeCenter HS20 Xeon 3.06 GHz, Gig-Ethernet / 280 IBM	**1051.4**	**1713.6**	Bayer CropScience GmbH/2004	Industry *Life Science*
360.	SuperDome 875 MHz/HyperPlex / 512 HP	**1024.3**	**1792**	BMW AG/2004	Industry *Automotive*
361.	SuperDome 875 MHz/HyperPlex / 512 HP	**1024.3**	**1792**	BMW AG/2004	Industry *Automotive*
362.	SuperDome 875 MHz/HyperPlex / 512 HP	**1024.3**	**1792**	BMW AG/2004	Industry *Automotive*
373.	DL360G3, Pentium4 Xeon 3.2 GHz, Myrinet / 256 HP	**1015.5**	**1638.4**	Siemens/2004	Industry *Electronics*
374.	DL360G3, Pentium4 Xeon 3.2 GHz, Myrinet / 256 HP	**1015.5**	**1638.4**	Siemens/2004	Industry *Electronics*
385.	SuperDome 875 MHz/HyperPlex / 512 HP	**1013**	**1792**	Magirus International/2004	Industry *Database*
409.	eServer Opteron 2.0 GHZ, Gig-E / 432 IBM	**987.1**	**1728**	FinanzIT/2004	Industry *Database*
412.	Integrity Superdome, 1.5 GHz, HPlex / 192 IBM	**971.2**	**1152**	BMW AG/2004	Industry
413.	Integrity Superdome, 1.5 GHz, HPlex / 192 IBM	**971.2**	**1152**	BMW AG/2004	Industry
414.	Integrity Superdome, 1.5 GHz, HPlex / 192 IBM	**971.2**	**1152**	BMW AG/2004	Industry
470.	SuperDome 875 MHz/HyperPlex / 448 HP	**897**	**1568**	Magirus International/2004	Industry *Database*
474.	Fire 15k/6800 Cluster / 672 Sun	**891.4**	**1209.6**	Aachen University - RWTH/2003	Academic
475.	Opteron 2.0 GHz, GigE / 360 NEC	**891**	**1440**	VW (Volkswagen AG)/2004	Industry *Automotive*
499.	SuperDome 875 MHz/HyperPlex / 416 HP	**850.6**	**1456**	Magirus International/ 2004	Industry *Database*

E. Glossary

Bandwidth: The maximal amount of data which can be transported in a time interval by a connection (memory buses, networks, I/O channels etc.)

Capability computing: One or, in many cases, several simultaneous demands on computing resources (computing time or power, central memory, I/O demand etc.) by a program for the solution of a problem are so high that they can be processed as a whole only by specially designed supercomputers

Capacity computing: Large numbers of (single) calculations (e.g., a parameter study which is supposed to cover a large range of parameters), in which the single program does not need especially huge computing resources but the number of program executions makes extremely high demands on turnover and computing performance

CPU: Central Processing Unit

Femto: prefix of a unit, meaning 10^{-15} units

Flops: floating point operations per second (measure for the performance of computers)

Giga: prefix of a unit, meaning 10^{9} units

Grid: System for the provision of data, information and computing services, in which the technical details of the concrete, typically geographically distributed realisation is hidden behind interfaces. The service is generally provided comprehensively with regard to system, location and organisation. The user, who is customarily member of a community (group with shared interests), sees only the “socket” from which he obtains the desired service but not the technical details which are necessary for the production of the computing, data or application activity. The name “grid“ is derived from “power grid“.

Latency: The time which elapses between the request and the reception of the first data bits.

LINPACK Benchmark: Standard benchmark in high performance technical computing which measures the computing performance during the numerical solution of very large linear equation systems. Rather obsolete for real-world application profiles because it correlates strongly with the peak performance of the computer but hardly or

not at all with other specifications such as bandwidth. LINPACK is the basis for the so called TOP500-List.

Nano: prefix of a unit, meaning 10^{-9} units

Peta: prefix of a unit, meaning 10^{15} units

Simulation: execution of experiments on a theoretical model, which can often be completely realised on a computer

Sustained Performance: Average performance of a computer in permanent operation. The kind of application to which the permanent operation refers, may vary (e.g., sustained benchmark performance, sustained performance for arbitrary applications, sustained performance for a representative cross section of applications etc.)

Tera: prefix of a unit, meaning 10^{12} units

TOP 500-List: List of the most powerful computers installed worldwide (www.top500.org)