

Teraflops For Europe

In little more than a decade, simulation using high performance computers has become an essential tool for theoretical physics. It is capable of solving a vast range of crucial problems inaccessible to conventional analytic methods. Computer simulation has effectively become the calculus for interacting many-body systems so it is the key to the study of transitions from isolated to collective behaviour.

The simulation of lattice gauge theories in high energy physics has acted as an ice-breaker in this development, clearing the path for many other applications. However, work on gauge theory demonstrates that we are still orders of magnitude removed from the computing requirements for physical solutions to quantum chromodynamics (QCD). Only a large jump to much more powerful computers will bring us to the next level of theoretical understanding.

A similar conclusion has been reached in condensed matter physics, quantum chemistry fluid flow studies and biophysics (see front cover illustration). The German government, for example, is spending several million DM to support supercomputer research in areas where many problems require a teraflop machine running for a day or longer. Some applications in meteorology (e.g. turbulence) need 1000 years.

US Response

At the beginning of October 1990, a collaboration of 16 leading North American universities and research laboratories, together with the Thinking Machines Corp. as an industrial partner, submitted a proposal to the US Department of Energy for the construction of a QCD Teraflop Supercomputer — the first computer that would be capable of providing a sustained performance of 10^{12} Tflop. Such a facility would make possible both the quantitative prediction of many experimentally crucial parameters as well as the search for new phenomena in non-perturbative strong interaction physics. The targets are one teraflop on a 32 bit word with 300 Gbyte of memory, 2 Tbyte of disc memory accessible at 3 Gbyte/s and 85% reliability with fault tolerance. The development calls for \$US 39 M in government funding to buy a stripped down version within three years. The possible existence of this general purpose machine, which would be commercially available in 1995 at the earliest, may set the boundary conditions for Europe's physicists.

Two other US companies (but not IBM) gave indications of future plans: Intel hopes to produce a 1 Tflop machine in 1995-96; Cray Research's architecture is to be exploited to 100 Gflop when a massively parallel configuration will take over.

European Situation

More than 50 European lattice gauge theorists, together with some experts from condensed matter physics and from computer science, met at CERN on 27/28 November 1990 with representatives from the

computer industry to examine if and how the decisive next step could be achieved in Europe. The situation is undoubtedly more complicated than in the USA.

There is ample expertise at the research level: the APE group in Rome, along with the Columbia University group in the USA, pioneered the construction of high performance, special purpose supercomputers for QCD. However, while 100 Gflop may be attainable at APE, an increase to 1 Tflop seems beyond the project's resources. A collaboration between groups from German universities, working at HLRZ, the High Performance Computer Centre in Jülich, has made significant progress in the study of QCD thermodynamics and hadron properties. Determined efforts by a group at Edinburgh University have recently led to the establishment of a UK Grand Challenge Collaboration for the study of non-perturbative QCD, albeit with only one-third of the requested funding (about £1 M).

The European counterpart to the American supercomputer industry however needs urgent and rapid development: companies such as AMT, Meiko and Suprenum must grow and enhance their products. One proposal is a European Community supported, Meiko — Suprenum joint venture — the general purpose GENESIS-P parallel supercomputer — to achieve by mid-1992 some 60 Mflop on a single node with the possibility of 2 to 4096 node configurations. In continuing its normal product development,

Meiko indicated that it hopes to offer a 250 Mflop processor by 1992-93.

Initiatives

The European lattice gauge theory community strongly supports a push into the teraflop regime and has endorsed further action towards this aim. An exchange of ideas and expertise could be extremely useful in developing the next generation of machines. The steering committee of the European Teraflop Meeting has therefore set up a Technical Working Group to assess the possibility of finding suitable partners from the European computer industry. The idea is that the European scientific community, which has expertise in the construction of large scale special purpose computers and in large scale computer simulation, would collaborate with such a partner. Given sufficient funding, this could lead to the construction of the first European Teraflop Supercomputer on a time-scale competitive with that of the USA. If successful, this ambitious enterprise would stimulate computer simulation as well as supercomputer construction in Europe.

Collaboration with other research areas, will also be actively sought. As there is unlikely to be a single funding body in Europe (the USA can count upon the National Science Foundation) one strategy may be to organize the supercomputer community so that financing can be obtained from both central (e.g. EC) and national sources.

H. Satz, CERN, Geneva

Towards A Supercomputing Initiative

Last month's report of a distinguished working group charged by the European Commission to recommend a European policy for supercomputing stressed that the field was on the verge of a new revolution. Scientists from several disciplines described how a new philosophy, half-way between theory and experiment, of conceiving science is emerging. Costly experiments will be replaced to a large extent by computer models that reveal "the finest details beyond the limits of any real experiment". Scientific competitiveness will then be governed by the availability of adequate computing resources, specifically a teraflop machine within 7-10 years, high speed data networks with 1000 times today's capacity, and innovative software to handle interactive real-time use as well as revolutionary changes in computer architecture. Also included in the five proposals for an integrated action plan involving an annual investment on the order of 1000 MECU over 10 years was the strengthening of education, research and development to exploit the new opportunities.

Prof. Carlo Rubbia, the Group's Chairman and Director-General of CERN, recognized at the press conference in Brussels to announce the proposals that the report established the "fundamental strategic importance of the high speed computing sector". Priorities and funding were another matter. Professor F.M. Pandolfi, the CEC Vice-President, indicated that a paper due in March could outline the strategic priorities (supercomputing is one) that may be covered by the second allocation of the Community's Framework Programme for R. and D. He emphasized, however, that the overall policy calls for partnership. For example, he agreed with the Group's recommendation that "technology transfer mechanisms" should be exploited to help strengthen the European software industry.

Europe with 30% of the market for high-performance computing is almost totally dependent on foreign suppliers. The Group felt that we stand on the threshold of a new era where competition with the USA and Japan becomes feasible with the development of parallel computing architectures using "non-exotic" hardware technologies (see above). The US Congress is meanwhile considering the Federal High Performance Computing Program involving a further injection of 2000 M\$ over five years for the installation of a huge fibre optic data superhighway and with 650 M\$ each for software and hardware development.