Overcoming the Isolation

Berni Alder, Professor Emeritus at the University of California, Berkeley, who pioneered the use of particle methods in hydrodynamics, took the opportunity as a keynote speaker at PC'94 to reflect on why computational physics, in spite of its successes, remains somewhat isolated in academic circles.

Although there remains the problem of handling the interaction between particles in simulating, for example, Schrödinger's equation, Monte Carlo methods represent a great success for computational physics. However, other major challenges have not yet been tackled successfully. They include three-dimensional turbulence where there has been little progress for 60 years as distance and time-scales are very large, and pattern recognition where the hopes for massively parallel computation have not materialized. There will be little progress in pattern recognition until much more complexity is built into computing nodes, thus mimicking neural networks. The folding of DNA molecules and why it takes place so quickly are poorly understood; a new paradigm is needed as random folding clearly does not work. People are groping for ideas, although it seems that the presence of water is crucial in helping guide the DNA molecules through metastable intermediate states (see figure).

What is perhaps sad is the fact that computational physicists are often not accepted as equal partners in the academic world in spite of their considerable achievements which include, for example, calculating the mass of the proton from first principles (the challenge now is to calculate the mass of the hydrogen atom). They have tended to seek homes outside mainstream physics in areas such as mechanical engineering and applied materials research. The situation arises because an undergraduate curriculum in computational physics has not been defined, and because there is some confusion as to what a computational physicist represents.

Claudio Rebbi, in summarizing below how a curriculum can be developed, goes a long way to explaining what computational physics tries to achieve. However, the community itself is partly responsible as it has not focussed sufficiently on theoretical aspects and the development of algorithms, but tended instead to be absorbed by advances in computer technology. Hence, in spite of considerable discussion and development, few qualitatively new algorithms have emerged in the field of massively parallel computing.

However, one should not be too critical as there are efforts under way to identify new paradigms. Consider, for instance, the field of computational fluid dynamics. Lattice-gas methods have not been successful in simulations as the largest computers can only handle 10^6 particles for 10^{-8} seconds. Monte Carlo methods have been shown to be very effective, but only at low Reynolds numbers. Using a hybrid approach that combines the simulation of particle scattering at boundaries with continuum hydrodynamic methods away from the interface, it is becoming possible to tackle more relevant high-Reynolds number problems. The crucial issue is to match the Navier-Stoke's formalism with particle-like simulations by ensuring that conservation laws are satisfied. The first results for a magnetic read-write head floating just above a rapidly spinning magnetic recording disc show that one can realistically simulate the development of fluid-flow instabilities. If Rebbi's ideas are widely adopted it should be possible to develop other paradigms, and envisage further successes.

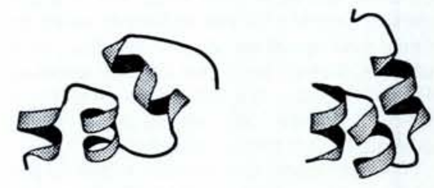

Predicting three-dimensional protein structures. Attempts to reliably predict the structures of proteins are in their infancy. One approach is based on the idea that native protein structures are in the state with the lowest free energy. The interaction between the protein and its water solvent is known to play a dominant role in the stability of the protein. This interaction can be introduced by assuming that it is equal to the sum over all atoms of the atomic solvation parameter times the accessible surface area. Braun and Mumenthaler reported at PC'94 that adding this simplified semi-empirical energy surface term to the standard conformational energy term and carrying out a minimization using Monte Carlo simulations gave folded structures with the correct three-dimensional topology. This is illustrated in the figure for a small test protein having three packed helices. One of the low-energy simulated structures (the structure on the left) closely resembles the structure on the right determined by NMR. So its seems that efforts to predict protein structures are moving in the right direction.

Teaching Computational Physics

Claudio Rebbi of Boston University, in arguing that the teaching of computational physics must become an integral part of physics education, offers some general principles as a guide.

COURSE PROJECTS IN COMPUTATIONAL PHYSICS

Calculate the thermodynamic properties of crystals.

Solve an equation for the internal energy of a crystal based on a quantized oscillator and the Debye approximation for the frequency of the normal modes of vibration. This introduces numerical integration. Going beyond the Debye approximation to use a more realistic dispersion formula demonstrates that the computational as opposed to the analytical perspective simplifies the calculation.

Calculate the trajectories of N-bodies subjected to mutual gravitational interaction.

A prototype project for the general class of molecular dynamic simulations. Introduces various algorithms for solving ordinary differential equations with initial value data. Leads to a discussion of stability and the need to employ more sophisticated techniques, such as the adaptive time step. With careful choices of initial data one can explore chaotic behaviour. Also introduces advanced computing technology (*e.g.*, implementation on massively parallel computers).

Calculate the energy of a superconducting vortex in the theory of Ginzburg, Landau and Abrikosov.

Minimize a function of the several variables that parameterize the field configuration. Introduces fundamental algorithms (*e.g.*, method of conjugate gradients).

Calculate the energy levels of a bound state of a heavy quark and antiquark with a potential that has a Coulombic component and a linear component.

Introduces boundary value ordinary differential equations. Goes beyond learning the algorithms to checking the predictions of the model against experimental data.

The majority of applications of computational physics involve numerical computations. The discipline must thus be taught by using a fair amount of numerical analysis. But there is much more to computational physics than the study of algorithms and numerical methods. Computational physics is the art of formulating and solving physics problems on the computer. It is this art which we must convey to students. To accomplish this we must keep the teaching of computational physics anchored to the treatment of actual physics problems. These can provide the background and motivation for the introduction of algorithms. The consideration of a concrete problem may, moreover, give the opportunity to illustrate the difference between an analytically minded approach and a computational approach, and how by changing one's perspective one can frequently achieve a simpler and more efficient solution. Finally, as code is developed for a teaching project, the students can refine programming skills and be exposed to advanced computing technology.

A Matrix

We therefore essentially have a framework for teaching computational physics consisting of a matrix of physics problems, numerical methods and computing technology. Such a framework provides an excellent educational platform, while also reflecting the integration of knowledge that plays a fundamental role in progressing algorithmic research. For we often find that developing more efficient ways to carry out some computation rests on understanding the physics of the phenomenon under study. In my own field of research, one of the most demanding computations consists in the calculation of quark propagators in the background of a non-uniform gauge field, which must be carried out many times to simulate quantum fluctuations. In mathematical terms, the calculation consists in solving a very large system of linear equations (with a number of variables that can reach into the millions), with a matrix of coefficients which is very sparse, but non-uniform and with poor-quality condition numbers. Ideally, one would like to apply multigrid methods or other acceleration techniques. What one sees is that the ability to implement such methods is closely coupled to understanding the physical properties of the lowest eigenstates of the quark field, so algorithmic progress and progress in unravelling the physics go hand in hand. Issues of physics and computing are not separate, but should be seen as complementary.

Putting into Practice

There is an emerging consensus that computational physics should be taught at the undergraduate level. This is a difficult challenge. An undergraduate course on advanced computing in physics based on the matrix approach was given for the first time last year at Boston University as part of a coordinated set of new undergraduate courses on parallel computing supported by the US National Science Foundation. Some of the projects used in the course are summarised in the insert.

My experience in teaching the course was that I assumed too much about the students' background preparation. This is a real difficulty because teaching computational physics must proceed concurrently with the rest of the curriculum. One way to help solve the problem is to use the tools whose utilization we wish to teach as the basis for making many mathematical concepts accessible. For instance, the notion of a Laplacian operator can be introduced using image manipulation procedures. The idea is to start with pixels that form an image and blur the image by replacing their numerical values with the averages of their neighbours.

A more serious — and I believe general — problem is to make room for new courses in an already overcrowded curriculum. Exposing students in general courses to computational projects is not a solution because there is much more to the methodology of computational physics than seeing the results of its applications. However, by cleverly integrating the power provided by modern computational tools in the overall curriculum we may be able to help students acquire more easily all of the many notions that are indispensable for a physicist.

This text is taken from an invited talk given by Claudio Rebbi at the *6th Joint EPS-APS International Conference on Physics Computing* (Lugano; 22-26 August 1994) and published in the proceedings.

PC'96: 7th JOINT EPS-APS INTERNATIONAL CONFERENCE ON PHYSICS COMPUTING

Pittsburg, PA, USA June 1996
Contact: D. Barnes, LANL

1997 IUPAP Computational Physics Conference

The International Union of Pure and Applied Physics (IUPAP), in reappraising its role, decided at its General Assembly last year in Japan to explore the possibility of forming a Special Commission for Computational Physics. A working group was set up which will shortly present a proposal to IUPAP's Executive Committee for the new Commission (the Executive has the authority to establish the Commission).

It is proposed that the Commission should adopt a broad perspective by promoting "the exchange of information and views among members of the international community of physicists in the area of of computational studies of problems related to physics including:
- numerical and symbolic models and algorithms for simulating physical systems;
- computational control and data processing in experiments;
- computing environments;
- the physical basis of computing machinery."

David Andersen from California who chairs IUPAP's computational physics working groups.

The new Commission also sees itself as the body that sanctions international conferences in the field of computational physics. The working group has therefore agreed with representatives of the various regional physical societies that a IUPAP-sponsored conference should be held every three years, with the joint EPS-APS conference alternating during the intervening years. Consequently, the next joint EPS-APS conference (Physics Computing '95) will be held next in Pittsburg, followed by PC'96 in Cracow with the IUPAP event in 1997 in Beijing.

SUPERCOMPUTING CENTRES

A Variety of Approaches

Supercomputing facilities in the various east and central Europe are evolving at different rates in the various countries, with Poland adopting the most adventuresome attitude.

POLAND

Poland's main supercomputing centre is the Academic Computing Centre based in Cracow's Institute of Computer Science. It was established some 20 years ago and now serves 11 institutes in Cracow, but following a recent policy change it will see its user community grow to include other institutes. Most investments have been funded by the State Committee for Scientific Research *(KBN)* which runs the country's research grants scheme (see *EN*, April 1994). The most recent investment is an agreement signed a few months ago to purchase one of Convex's latest generation of massively parallel supercomputers (a 16-processor machine in the Exemplar series that offers scalable parallel processing with from 1 to 128 central processing units). The Centre already has three Convex machines, the first of which — and the first Convex in central and eastern Europe — was acquired about four years ago. Some six other Convex computers are distributed around the country (in Torun, Warsaw University, the Oil and Gas Institute in Cracow, the Institute of Nuclear Physics) and two are in industry.

Warsaw's Centre for Scientific Supercomputing based at the Institute of Applied Mathematics, Computer Science and Mechanics of the Polish Academy of Sciences is equipped with a so-called "baby" Cray (an 8-processor machine in the Cray Y-MP EL series with up to 32 CPUs). Warsaw University also has a computing centre based on an IBM 3090 which was delivered in 1990 to form an European Academic Research Network (EARN) node — the first in east and central Europe — as part of IBM's Central European Academic Initiative. There is a computer centre attached to the Polish Academy of Sciences Institute of Biology in Poznan which acquired a 4-processor Cray Y-MP EL this year. The *KBN* recently signed an agreement with Cray for 6.5 M$US to be spent over 2 years to upgrade supercomputing resources.

CZECH REPUBLIC

The Institute of Physics of the Czech Republic's Academy of Sciences acquired a Cray Y-MP EL in 1992. Owing to financial constraints, it has proved difficult to develop a large user base (users are concentrated in physics, with most coming from universities). The Institute recently formed a consortium to seek support from the Czech government's Foundation for the Development of Science to build up a National Centre for Supercomputing that would be in a position enlarge its client base and support by working for industrial companies. These could include US firms that have been contracted by the government to upgrade Soviet-designed nuclear reactors. The immediate need is to increase the number of workstations and improve software resources. However, a project along these lines was not approved this

A deskside Y-MP EL Cray supercomputer.