WORKING DRAFT - NOT FOR PUBLICATION

ever, I will discuss below an approach to solving partial differential equations which involves few standard arithmetic operations, and should be implemented very efficiently on the Connection Machine. The approach is based on cellular automata, and can be considered as intermediate between a simulation of molecular dynamics and the standard continuous partial differential equations. Essentially the information say on fluid density that is usually given as a real nomber at a single lattice site is represented by an average of the values of several sites each with a discrete single lattice site is represented by an average of the values of several sites each with a discrete single lattice site is needed to analysis is needed to analyse these est of possible values. A new, statistical, form of numerical analysis is needed to analyse these algorithms. I will describe the progress on this so far below.

equations on 256x 056x 256-site stids. The physics is such that fuid motion on very

Scientific Computation with the Connection Machine

Stephen Wolfram

(March 1985; revised April 1985)

1. Introduction

The biggest application of today's most powerful general-purpose computers is for scientific and mathematical computations. These notes discuss how it should be possible to do many of these computations much better with the Connection Machine.

The point of most scientific computations is to work out how physical systems will behave. Usually the computations correspond essentially to explicit simulations of the behaviour of the physical systems. (There are fundamental reasons to expect that the evolution of physical systems is often "computationally irreducible", so that such explicit simulation is the only possible procedure.) The structure of the computations required then depends on the structure of the physical system studied. Now most physical systems of interest can be considered to consist of a large number of identical components, such as molecules in a gas. When such systems are simulated on standard serial-processing computers, each component must be treated in turn. But a parallelprocessing computer such as the Connection Machine can process information in a way much closer to that of the physical system itself, and can treat computations for many components in parallel.

Another feature of physical systems is that the interactions between components tend to be local; as a consequence, the computations required are much like those in a cellular automaton they can be implemented on a finite-dimensional lattice, without additional long-range connections. As a consequence, it should be possible to write programs for them without addressing the problems of parallel processing in full generality.

The vast majority of large-scale scientific computations involve the solution of partial differential equations. Most often, these equations are for the behaviour of some form of fluid. Although fluids are ultimately made from discrete molecules, the equations approximate them by continuous macroscopic variables. Whenever the equations are simulated on a digital computer, some form of discrete approximation must however be made. Conventional numerical analysis takes a discrete grid in space and time, but leaves the variables at each site as high-precision floating point numbers. Conventional algorithms for solving partial differential equations then involve extensive arithmetic with these floating point numbers. Since the Connection Machine has no specific hardware for such arithmetic, it is comparatively inefficient for such algorithms. However, I will discuss below an approach to solving partial differential equations which involves few standard arithmetic operations, and should be implemented very efficiently on the Connection Machine. The approach is based on cellular automata, and can be considered as intermediate between a simulation of molecular dynamics and the standard continuous partial differential equations. Essentially the information say on fluid density that is usually given as a real number at a single lattice site is represented by an average of the values of several sites each with a discrete set of possible values. A new, statistical, form of numerical analysis is needed to analyse these algorithms. I will describe the progress on this so far below.

2. Frontier scientific computations

Here is a list of some current significant basic scientific research problems that the Connection Machine should be able to make progress on.

A. Fluid turbulence

Complicated irregular, or turbulent, flow is very common, and is important for many technological applications such as aerodynamics. But it is still not known exactly how or why turbulence occurs. It has not been possible to solve the Navier-Stokes partial differential equations in systems with realistic fully-developed turbulence. The best current simulations approximate the equations on $256 \times 256 \times 256$ -site grids. The physics is such that fluid motion on very large scales is essentially determined by overall boundary conditions, and on very small scales is smoothed out by viscous effects. Current simulations really do not have enough grid points to see an intermediate ("inertial") range of scales, on which true turbulence can occur. And in fact it is not known what the fundamental mathematical mechanism for turbulence is. The nature of the mechanism is important in determining the kinds of approximations that can be made in solving the fluid equations, while still correctly reproducing turbulence.

One explanation for turbulence is that it results from amplification of external noise or perturbations on the system. There is evidence that this is part of the story, but the fact that experiments in liquid helium have yielded almost exactly repeatable complicated patterns of flow shows that it is not the whole story. The second explanation, now popular through studies of strange attractors and so-called "deterministic chaos", is that the flow is arbitrarily sensitive to initial conditions. Given a real number as the initial data, the pattern of flow progressively "excavates" more and more digits in this real number. So to reproduce turbulence one would need to work with arbitrary precision real numbers. In addition, with either of these two explanations, one could not in practice expect to find the form of the flow in any particular case, because it would be determined by unknown small details of initial conditions. In fact, it is neither necessary nor probably necessary to know the exact, irregular, flow. What is needed is to know statistical averages that capture what regularities do exist in the flow.

I have recently come up with a third explanation for turbulence (that concurrs with intuition from many numerical fluid experiments, but could not be formalized before). I think that fluids can act much like pseudorandom number generators, producing complicated and apparently random patterns of flow even from simple initial conditions, analogous to the simple "seeds" of pseudorandom number generators, or simple "keys" for cryptographic systems. This phenomenon is very common in cellular automata, and there is every reason to expect that it is also present in real fluids. And with this explanation, it suffices to use discrete approximations to reproduce the fundamental features of turbulence.

I am not sure exactly what the best system to simulate turbulence in is. The Taylor-Green vortex has been a popular one, but it may be too special for standard fully-developed turbulence to occur. However, it might be a good one to start on, not least because new results could be checked against old ones. There are several specific questions of interest. First, can classical fully-developed turbulence occur in the simulations at all? (Perhaps the Navier-Stokes equations just break down in the turbulent regime, and one must use explicit molecular dynamics.) Do the Kolmogorov dimensional scaling laws work? And what corrections, due for example to intermittency, occur? Then one can ask about the sensitivity of the flow to perturbations. Do all perturbations grow exponentially in size (so that they have positive Lyapunov exponents)? Or are (as I suspect) sufficiently small perturbations effectively damped, at least if one looks at features of the flow above some small length scale? In addition, what level of discrete approximation suffices to reproduce what features of the flow? The coarser the approximation that nevertheless reproduces relevant features, the more useful it will be for practical engineering purposes. Finally, does the flow really behave like a pseudorandom number generator, as I claim? Are the statistics of flows with very simple initial conditions indistinguishable from those with more complicated initial conditions?

B. Lattice spin systems

Lattice spin systems are used to model a variety of statistical mechanics systems such as ferromagnets. Scale invariance implies that their properties at least at certain phase transitions are universal across many different systems. Some properties of some of these universality classes have by now been well studied, mostly by computer simulation (in only a very few cases are exact analytical solutions known). But computer simulations have so far been inadequate to investigate other classes. In addition, typically only very gross statistical properties have been studied. There has been rather little study of either time-dependent (dynamic) phenomena, relating for example to the approach to thermodynamic equilibrium, or of spatial patterns such as domain wall structures.

The simplest spin systems of interest are Ising models, consisting of a lattice of identical spins, each with two possible states, and each interacting only with their nearest neighbours. An analytical solution is know for the two-dimensional Ising model. The three-dimensional Ising model has been much studied numerically, and it seems that only details remain to be found. But there are many generalizations of the Ising model that should be studied further. Some generalizations (e.g. Potts models) take more possible values for each site. Other consider longer range interactions. And still others make the system inhomogeneous, taking different randomly-chosen interaction strengths for couplings between spins at different positions. Such generalizations are known as spin glasses, and are found to have very rich behaviour.

One of the first questions about a spin system is whether it exhibits a phase transition. If so, what type of phase transition? And what is the relevant order parameter (macroscopic quantity analogous to magnetization for a ferromagnet that identifies different phases)?

To find the behaviour of a spin system in thermal equilibrium, one must sample its possible configurations, weighted with the appropriate probabilities for thermal equilibrium. A method that has been used extensively is to start from one configuration, then move to other configurations by making perturbations determined by a pseudorandom number generator, and thereby obtain a Monte Carlo sampling of configurations. A promising recently-proposed method instead uses deterministic cellular automaton rules to move from one configuration to another. (This procedure mimics the actual physics of spin relaxation in the case where interactions between spins dominate interactions with lattice vibrations; it effectively generates randomness through an autoplectically.) To get accurate results, one must sample as many configurations as possible, and these configurations must be as uncorrelated as possible. But even given the configurations there is also substantial computational work involved in computing statistical averages, for example of candidate order parameters, particularly when these are complicated nonlocal quantities that depend on the values of many spins.

The fastest code for computers like the Cray-1 updates about 10⁷ sites per second in a standard Ising model, using a Monte Carlo approach, and about 10⁸ sites per second using the cellular automaton approach. Special purpose shift-register-based computers have built for doing these simulations (one at Santa Barbara, one at AT&T Bell Labs); both run at about the same speed as the Cray-1. The Connection Machine should be able to do perhaps a factor of a thousand better.

Unfinished in 1985

There are some specific questions about spin systems that would be of particular interest. The most significant probably concern spin glasses, or certain imitations of them. One of these imitations is the random field Ising model - an Ising system with an external magnetic field randomly applied at each site. An important question, not yet answered definitively, is whether there is a phase transition in this model in two space dimensions. For spin glasses, it is important to try and find out what order parameters characterize the phase transitions which physical experiments suggest are present. And to try and determine the answer to related questions about the exponentially large number of metastable states that appear to be present, and lead to a very slow approach to thermal equilibrium.

I suspect in fact that it is in general formally undecidable whether a spin system exhibits a phase transitions, so that there is no upper bound on the amount of computation required for any particular case. And in fact it seems that much computation is often required. There are some partially analytical methods (relying for example on series expansions) known, but direct simulation is usually best for addressing fundamental questions of interest.

C. Lattice quantum field theories

Through the path integral formulation quantum field theories can be put in a form that is mathematically analogous to spin systems. Like partial differential equations, however, they are intrinsically continuous systems, but must be approximated for digital simulation by discrete four-dimensional lattices. The most interesting case for study is quantum chromodynamics (QCD), the theory of quarks and gluons. A fundamental question is whether this theory implies absolute confinement of quarks and gluons into the observed hadrons (such as protons and pions). Such confinement is related to the absence of a phase transition between the regime of strong coupling covered by certain approximations, and the weaker coupling regime realized in practice. There is strong experimental evidence for confinement, and some theoretical evidence. Calculations with QCD should in principle allow properties of particles such as protons to be calculated theoretically. But in practice simulations have used lattices which are too small for very useful results to have been obtained (typically at most 16⁴ sites).

There are two very different kinds of computations done in studying QCD: those with quarks (or other fermions), and those just with gluons (which as bosons). When only gluons are considered, the quantum field configurations can be sampled much as in a spin system. In this case, however, the "spins" are continuous variables (representing the "direction" of the gluon field at each point in the internal gauge group space). But these may be well approximated by a discrete set of possible values - 4 or 5 bits per site seems to suffice. (In fact, the possible values are best taken to be elements of finite subgroups of the continous gauge group, or elements of products of such subgroups.) This discrete approximation is adequate in the strong and intermediate coupling regions of most practical interest; but in the weak coupling region where contact might be made with perturbation theory results, the site values are close to one, and more accuracy is needed. (If only a few site values are allowed, fake freezing phase transitions are found.)

When quarks (or other fermions) are included, things become much more complicated. There are global constraints on the possible configurations, resulting roughly from the Pauli exclusion principle, which applies to fermions, but not bosons. In sampling configurations one must find inverses of large matrices. Nevertheless, at least in one approximation, these inverses can be found by calculating the Green's function for the Dirac differential operator using an iterative local relaxation method. However, it seems that higher numerical accuracy must be maintained in such calculations. Of course, one must work with quarks if one is to find out the properties of particles such as protons or pions which contain them.

There is considerable theoretical interest in further studies of the pure gluon system. One would like to have a better qualitative picture of the structure of the ground state for the system. A few gross properties of it have been investigated; but the overall strucutre remains the subject of speculation. (In fact, a better understanding may require not so much extensive lattice gauge theory computation as sophisticated graphics to visualize the results.) Then there is the question of excitations, which represent particle states ("glueballs"). Finding for example the masses of

Unfinished in 1985

such objects requires measuring exponentially-decaying correlations between different parts of the lattice. It seems that at least 24 bit numerical accuracy is needed in forming these averages for useful results to be obtained. In current simulations, the computer time spent finding averages is comparable to that spent in updating the spins ($O(10^4)$ sweeps are typically done in total); to obtain better statistics it is important to make more sweeps over the lattice to remove correlations between successive configurations sampled, which give non-Gaussian errors.

In addition to QCD, a lattice quantum field theory that is now attracting attention is a version of quantum gravity based on discrete triangulations of spacetime called Regge skeletons. Here the problem is to enumerate a large number of objects best represented as graphs; the Connection Machine, using long-range connections, may be very suitable for this. First one must check that weak field approximation phenomena, such as gravitational waves, can be recovered in the model. Then one should be able to tackle a fundamental current problem: why the effective energy density in the universe as it affects gravitation is so small, even though zero point fluctuations in quantum fields should make it very large (or infinite). To address this question, one should compute the expectation value of the Ricci scalar curvature as a function of temperature (corresponding to inverse periodicity in the imaginary time direction) or matter density in a model universe. These necessary calculations will undoubtedly be most lengthy, but there is probably still some more exploratory work to be done.

D. n-body problems

This is a class of problems that concern the motion of bodies, usually approximated as point particles, interacting through various force laws. They involve the solution of sets of coupled ordinary differential equations, with typically one equation for each particle. Parallel processing should be worthwhile if there are many particles. However, current methods require solving the equations using standard numerical analysis in terms of high-precision (though usually fixed point) numbers.

Classic n-body problems concern the motion of several bodies interacting through graviational forces. Thus for example a problem that remains of current interest is to simulate the evolution of the solar system, and to determine for example whether it is stable. Related problems concern calculations for example of asteroid motion, which is in some cases expected to be chaotic, rather than regular and periodic. In the past, the most effective methods for celestial mechanics used symbolic methods in terms of the Fourier transforms (Poisson series) of orbits; but it seems that direct simulations are now becoming the better method.

Another class of few body problems concern the structure of medium-sized molecules, and the dynamics of reactions between them. These problems involve solution of the Schrodinger equation. In finding molecular structures, eigenvalues of large matrices must be found (with matrix elements representing weights of different basis states in a variational wave function). In studying reactions, one solves many coupled ordinary differential equations for complex wave functions.

An important class of problems concern the statistical, aggregate, behaviour of collections of very large numbers of particles. Two rather different cases occur: those in which the forces are of short range so that particles can be considered to interact through collisions involving a small number of bodies, and those in which the forces have a long range, so that many particles have a continuing effect on the motion of each particle. The forces between molecules, say in a liquid, are short range. Those due to gravity between stars in a galaxy are long range. Those between ions or electrons in a plasma are fundamentally long range, but the screening of charges in such systems makes them effectively short range for most practical purposes. In all n-body problems, whether short or long range forces are involved, the fundamental problem is to simulate the motions of as many particles as possible.

Probably the most important current n-body problem with short range forces concerns the structure of liquids. There are theoretical suspicions that molecules in liquids are often ordered, something like in a crystalline solid, but following an irregular lattice, perhaps based on

Unfinished in 1985

tesselations of icosahedra. I think there is a need for more extensive simulations to study these questions (which are often investigated by analogue models, such as bags of plasticene balls), though many of them may be quite feasible with standard computers.

But it is the gravitational n-body problem that is the most challenging. What is the behaviour of a self-gravitating gas, for example of stars in a galaxy or globular cluster, or of density perturbations during galaxy formation in the early universe? Unlike a conventional gas that follows the Second Law of thermodynamics, self-gravitating gases can start from disordered states and end up clumped in rather ordered ways. How this happens is just not known. There are some linear analytical approximations, but the full problem must be investigated by numerical simulation. And existing simulations have not had enough particles, or been done for enough time, to see any form of equilibrium attained. The calculations are difficult, and it is not clear to what extent high precision numbers must be used in them. An *n* particle with long-range interactions could require $O(n^2)$ calculations. (The naivest blocking procedures nevertheless apparently give incorrect answers.) The general problem to find out the overall behaviour of self-gravitating gases. Does an analogue of the Second Law of thermodynamics hold? (One suspects that the usual Second Law does hold in some form, but with entropy carried by the gravitational field.) If so, what is the form of the equilibrium ensemble?

E. Cellular automata

Cellular automata seem to capture the mathematical essence of many complex natural systems, including several of those discussed above, and they are very directly amenable to efficient simulation on the Connection Machine.

We have made extensive studies of one-dimensional and to some extent two-dimensional cellular automata. But we have never had sufficient computer power to make a serious study of the three-dimensional ones, which are directly relevant to many physical systems (such as fluids). Beyond the computational power necessary to update the cellular automaton configurations, there is also the significant problem of displaying three dimensional images, that consist basically of arrays of solid blocks with different colours.

F. Number theory

The structure of computations in pure mathematics, unlike those in physics, are not determined by the structure of real physical systems, and so need have none of the homogeneity or locality of actual physical processes. A? a consequence, it is less clear that they can be implemented efficiently through parallel processing.

A prototypical number theoretical problem, of some interest for public-key cryptographic systems, involves factoring large integers into primes. There are a variety of algorithms known, with rather different structures, but all taking a time more than polynomial in the number of digits of the integer to be factorized. The best known algorithm takes $O(e^{\sqrt{\log_e n \log_e \log_e n}})$ time (counting only arithmetic operations). It involves several kinds of computations. First, arithmetic with integers containing many digits. Second, arithmetic with certain low precision (perhaps 16 bit) floating point approximation. And third, finding greatest common divisors. Quite a lot is known about implementing arithmetic operations in parallel. But gcd appears to have escaped parallelization (and may not even be in the complexity class NC). Apparently, the algorithm stores many integers, then seives through them looking for ones with small prime factors. Presumably this seive could be implemented in parallel. In fact, it seems that the MPP has been used to advantage with this algorithm. But at present only numbers with at most about 75 digits can be factored.

There are numbers of several special forms whose factorization might be of interest. The best known are Mersenne numbers, of the form $2^n -1$. In general, it is not clear that there is much fundamental interest in factoring numbers per se. But there is a \$100 prize on offer for the first group to factor a 125 decimal digit number.

3. Production scientific computations

4. Conclusions and recommendations