Three experimental modeling systems

Francis Muir

ABSTRACT

Three experimental dynamical models of computation, M1, M2 and M3, supplement the lattice gas as alternatives to schemes based on PDE's. These models share a common structure with the lattice gas, and may solve similar problems; they differ substantially in how they represent data, and in the sharply defined principle that controls collision rule selection.

M1, like the lattice gas, is a stochastic model. The model field is a set of realizations of a random process. Field elements are binary, discretized in time, space, and velocity, and are initially selected with probability depending on the data value. Unlike the lattice gas, space and time sampling is not necessarily fine, and the crude representation of the data in any one copy is compensated by the number of copies. Evolution proceeds by translation and collision; the latter, probabilistic and irreversible, based on an entropy principle.

The second model, M2, replaces M1's set of binary fields with an ensemble-average of such fields, and the field elements now range from zero to unity. At each step and at each point, these averages are disassembled into their component (M1-like) binary vectors, which are then replaced by constraint-equivalent averages, and re-assembled.

The third model, M3, specializes M2 for applications where the initial conditions are close to equilibrium; the problem is now linear, and the field consists of vectors of real numbers with the collision process replaced by an idempotent matrix operation on the velocity vectors.

While none of the three model systems is reversible, all have stable inverses. M2 and M3 preserve most of the attractive features of the lattice gas, and may generate very efficient and noise-free schemes, since they represent the data economically and have no random element.

INTRODUCTION

In a previous report (Muir,1986) I introduced the lattice gas and discussed modifications for propagation through a heterogeneous medium. In this report I discuss three related dynamical systems that have come out of my experimental work, but first I review motivation, and take some time to discuss the structure of the lattice gas, which is the jumping off point for my three model systems.

The lattice gas revisited

D'Humieres and his co-workers (1986), authors of one of the three seminal papers in the field, have described the lattice gas thus:

> Lattice gas models are discrete in space, time and velocity. They consist of particles, which may be viewed as Boolean molecules, residing at the sites of a regular lattice. Their discrete velocities are chosen such that in unit time they can propagate to a neighboring site. There may be several particles at a given site, but each velocity direction is subject to an exclusion principle of (at most) single occupancy. The updating of the lattice is done by alternating propagation and collision steps; the latter take place at each individual site and are chosen to have the same conservation laws as the true micro-world (particle number, momentum,...).

To this I would add that energy is also a conservable quantity, covered by mass conservation if the particles travel with only one speed.

The principle attraction of the lattice gas as a dynamical model for computation is the molecular analogy: particles move, collide, and move on, honoring whatever conservative laws are prescribed by the problem. They form closed systems, are not restricted to small departure from equilibrium, and are capable of modeling complex behavior, such as turbulent flow, in a natural and bounded manner and without recourse to explicit description.

From a computational point of view, the Boolean field variables, simple arithmetic, and local and synchronized demands of the evolutionary process are well suited to present day supercomputers, and particularly so to massively parallel designs.

Less certain, but in the end perhaps more appealing, is the notion that lattice gas and like systems may be more accessible to applied physicists than are finite difference schemes, with possibilities of a straightforward calculus that would allow direct model building without benefit of numerical analysts.

Symbols

For simple 2D regular lattices, particle configurations are simply illustrated by

an ordered binary representation, in curly braces, of the state—ones where particles are, and zeros where particles are not. It is useful to chose the state index as the integer which shares the same binary representation. So that:

> **{1,0,1,0,1,0}** is the (hexagonal lattice) configuration where, clockwise from a preferred direction, the first, third and fifth positions are occupied by particles, and the second, fourth, and sixth are not. The state index would be 42.

Classic schemes

Three key papers in the lattice gas literature discuss what are now three classic models. The first is HPP (Hardy, 1976), a 2D model based on a rectangular grid and up to four particles at each grid point, each pointing towards one of the four nearest neighbors. The collision rules for mass, momentum and energy conservation are:

$$\{1,0,1,0\} \text{ into } \{0,1,0,1\}$$

and

$$\{0,1,0,1\} \text{ into } \{1,0,1,0\}$$

and all other states remain the same.

This scheme is fine for the diffusion equation (where the rules are modified to eliminate the momentum constraint) but has viscosity anisotropy problems with the wave equation and is quite unsuitable for Navier-Stokes problems where the necessarily simple rules introduce a spurious line momentum conservation that is at odds with the requirement that orthogonal velocity terms be allowed to convect the momentum to produce eddy flow (Wolfram, 1986a, 1968b). This scheme was the basis for the work in my last report, where the anisotropy problems were masked by high noise level.

The second model is FHP (Frisch,1986); also 2D, it is a 6-particle, nearest-neighbor scheme based on a hexagonal grid. This scheme was introduced to correct the HPP deficiency, since the three-particle rules:

$$\{1,0,1,0,1,0\} \text{ into } \{0,1,0,1,0,1\}$$

and vice versa break up the spurious constraint. In addition, this scheme has the two-particle, zero-momentum rules:

$$\{1,0,0,1,0,0\} \text{ into } \{0,1,0,0,1,0\} \text{ or } \{0,0,1,0,0,1\}$$

and the two other like exchanges. Both HPP and FHP can be projected on to the line to give 1D schemes.

There is no simple 3D scheme that has the isotropy of the FHP, but, fortunately, there is such a 4D structure, HLF (d'Humieres,1986), that can be projected down on to the 3D rectangular lattice. This is a 24-particle scheme, based on the next-nearest neighbors of a point on a rectangular 4D lattice.

Collision Rules

The 4- and 6-particle HPP and FHP models exist in, respectively, 16 and 64 different states, and rule design reduces to global examination. In contrast, the 24-particle HLF scheme mentioned in the last section has over 16 million states, and it should not be a surprise that no algorithmic prescription has been offered in the literature. This size problem appears to call for some ingenuity, although a Monte Carlo solution is not inevitable, since a 48 Megabyte look-up table is feasible.

Efficiency and precision

Questions have been raised on the efficiency of lattice gases for non-linear fluid-dynamical problems, but there should be a greater concern for linear modeling, where signal/thermal-noise is directly related to the distance of the model state from equilibrium, and yet where linearity depends on a vanishingly small such distance. For a fixed amount of computation there is an inescapable trade-off between thermal noise and non-linearity. Lattice gases may be precisely repeatable in a formal sense, but there is no escaping system noise and non-linearity distortion, whose attenuation may be bought at high cost.

The new models

These concerns, and an interest in better understanding lattice gas behavior and design, are the driving force behind the three models that form the content of this paper.

M1 SCHEMES

Comparison with the lattice gas

As stated in the abstract, M1, like the lattice gas, is a stochastic model. Model values are Boolean, and have no independent meaning, only when taken in aggregate. With the lattice gas the model is a micro-model, finely sampled in space and time, and the macro-model only emerges on averaging the micro-model in space and time. With M1, the model field is a set of realizations of a random process, and useful results only emerge when the individual realizations are averaged over the set. The M1 process is less physical and more information theoretic.

A collision principle

The other major difference between the lattice gas and M1 is in the collision rules. Whereas lattice gas rules are experimental, dealer's choice, M1 rules are uniquely determined for any lattice and particle configuration by an entropy principle.

> *Particles emerge from a collision in any of the constraint-equivalent states with equal probability.*

Note particularly that this allows for the particles to emerge unchanged in configuration, as indeed they must where there are no other equivalent configurations. As an example, and apart from do nothing rules, the **HPP**-equivalent rules are:

$$\{\mathbf{1,0,1,0}\} \text{ into } \{\mathbf{0,1,0,1}\} \text{ or } \{\mathbf{1,0,1,0}\}$$

and

$$\{\mathbf{0,1,0,1}\} \text{ into } \{\mathbf{0,1,0,1}\} \text{ or } \{\mathbf{1,0,1,0}\}$$

both with equal probability.

Reversibility and Invertibility

An interesting and quite general feature of this rule design is that although it is not reversible, it is time symmetric—there is precisely the same information (or lack of it) about the outcome of a collision as there is about its antecedent. This symmetry, coupled with the memoryless property, classifies M1 as a symmetric Markhov process. Since translation, the other half of the evolutionary process, is reversible, we can state an inversion rule:

> *Inversion is accomplished by reversing the translation operation and leaving the forward collision rules intact.*

Computational considerations

From a computational point of view, 1024 independent but structurally identical problems may be more efficiently processed than the 32-fold increase in each of the two spatial dimensions of a lattice gas model experiment, and this might be true at both ends of the computer power spectrum. Small machines, such as the IBM-PC, with limited memory and addressing capability, should benefit from the serial opportunities afforded by the multi-independent structure, while large machines, such as the Cray XMP4, will take advantage of the parallel possibilities.

M2 SCHEMES

The finite set of experiments that characterizes M1 is replaced in M2 with the average over the ensemble of all such possible experiments. Model values are expectations, ranging in value from zero to unity, and since these expectations are also probabilities of occurrence of particles, M2 schemes are Markov random fields.

Working with ensemble averages rather than finite sets of individual realizations, thermal noise is eliminated, and there appear possibilities for massive improvement in computational efficiency over M1, particularly for the modern parallel, floating-point multi-processor.

The collision principle and a dilemma

The maximum entropy principle which guided evolution in M1 is just as valid in M2. Given an input vector:

$$P_{in} = (pin_1, \ldots, pin_6)$$

find an output vector:

$$P_{out} = (pout_1, \ldots, pout_6)$$

which maximizes:

$$Entropy(P_{out}) = -\sum_{i=1}^{6}(pout_i \cdot log_2(pout_i) + (1 - pout_i) \cdot log_2(1 - pout_i))$$

subject to the constraint that the mass and momentum of P_{in} are conserved in P_{out}. The problem is that, unlike M1, there appears to be no closed form solution to this constrained maximization.

A stosszahlansatz solution

This solution follows a three step process. It is not optimal, the output vector is not, typically, the maximum entropy vector, but it honors the model bounds, $0 \leq pout_i \leq 1$, and the conservation constraints, and *may, on repeated application, converge to the (unique) constrained maximum entropy value from any allowed input.* The steps are as follows:

1. The input vector is disassembled into the set of all possible states weighted by their probabilities.

2. Each weighted state is replaced by the weighted average of all the constraint-equivalent states.

3. These weighted average states are reassembled into the output vector.

The disassembly follows a non-committal (maximum entropy) linear decomposition, which can also be viewed as based on a model of molecular chaos, or independence of motion between particles on different direction planes. Relevant results from finite state probability theory are:

1. *The expectation or ensemble average of a Boolean random variable is equal to the probability that it is unity, and the probability that it is zero is one less its expectation.*

2. *The joint probability of occurrence of a set of independent events is the product of their individual probabilities.*

The first allows for the determination of the probabilities of occurrence or non-occurrence of the various particle types, and the second for the various particle type combinations or state probabilities from the individual type probabilities.

Another strategy

The makings of an alternative strategy, which might be appropriate to the 24-particle HLF model, is as follows:

1. Establish the constraints.

2. Fit the momentum vector to the minimum number of points necessary to describe it.

3. Spread the excess mass out evenly over all the other points less those antipodal to the fitting points.

Inversion

Inversion follows the M1 pattern closely, and for the same reasons:

> *Inversion is accomplished by reversing the translation operation and leaving the forward collision process intact.*

M3 SCHEMES

Evolution

The acoustic wave equation is a linear equation based on more general forms by considering small perturbations about an equilibrium state. In a similar fashion the M2 evolution operation can be linearized in case the model is close to equilibrium—this being particularly useful since the constraints, the three conserved quantities,

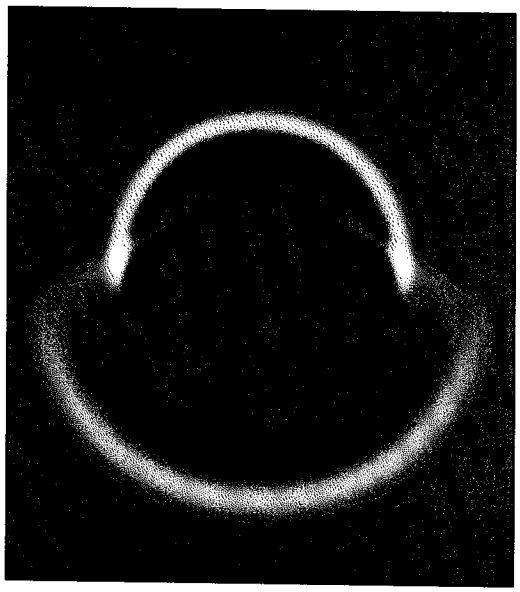

FIG. 1. A snapshot of a wave initiated by an impulsive point source in a low velocity medium. The wave has met a medium with twice the velocity but no change in impedance. The direct, refracting, reflecting, and head waves are all visible, and the reflected wave correctly has no amplitude at normal incidence. Notice also the isotropy in the direct wave, and the Hankel tail evident on both the direct and transmitted waves. Evolution took place every other time step in the slow velocity upper half as previously discussed (Muir, 1986). This model used M3, the linear model, a hexagonal grid, 6 particles, and the matrix operator discussed in the text.

mass, momentum and energy, are all linear functions of the M2 average model vector. I state without proof that the required conservation-constrained maximum entropy solution approaches the constrained L_2 solution as equilibrium is reached. From all this the matrix operator is found without difficulty, and in closed form. For example, in the case of the FHP-like M3 model, the symmetry in the placement of the six neighbors forces circular symmetry on the matrix. In turn, this fixes the eigenvectors of the matrix, and, by inspection, three of the eigenvalues must be unity to satisfy the mass and momentum constraints, and the remaining three must be zero to satisfy the minimum norm criterion. For this model the top row of the circulant is:

$$(1/2, 1/3, 0, -1/6, 0, 1/3)$$

and it was this operator that was used to generate the wave-field illustrated in Figure 1.

Inversion

Since the collision matrix is always idempotent, with either zero or unity eigenvalues, the appropriate inverse is the Moore-Penrose pseudo-inverse, the matrix itself. The inversion process thus follows M1 and M2 very closely:

> *Inversion is accomplished by reversing the translation operation and leaving the forward matrix operator intact.*

DISCUSSION

The uniform rules developed for the M1 systems could as easily be used, with profit, in lattice gases. It is likely that these rules could lead to much simplified analysis.

In case of linear diffusion and wave equation problems, the lattice gas is no match for M3, which is truly linear, noise free, and efficient, particularly on the super(mini)computers with floating-point processors that are the backbone of modern scientific conputing. Constant-q wave equations in their PDE form lie between the diffusion and wave equations, and an interesting possibility is that constant-q behavior can be mimicked by reducing the two momentum-controlled eigenvalues in the FHP-like M3 scheme from unity to something less.

The fine spatial sampling associated with the lattice gas also means that the mean free path, on the final macro scale, may be considerably shorter than a comparable M1 scheme, and this might be critical for some problems.

Lattice gases (and, I suppose M1 schemes) are adept at modeling eddy and turbulent flow representing metastable and not true equilibrium states. On the other hand, it is likely that M2 can reach equilibrium states that are not accessible to M1 and the lattice gas. Both representations may be useful, and it is possible that an equilibrium, M2 solution represents a useful point of departure for nonlaminar flow studies, since it is not difficult to switch back and forth between M2 and M1, and M2 and the lattice gas.

EQUIPMENT

Good hardware and software tools played an important role in the development of the ideas behind this paper.

Computer programs were developed on an IBM PC-AT, using a Logitech mouse, and the Norton Editor, Commander, and Utilities. Pilot code was compiled, linked, and executed with the Ryan-McFarland FORTRAN77 system. pfs:ACCESS was the communications program that provided the link over to a main-frame UNIX machine, the Convex C-1 of the Stanford Exploration Project. Remote operation of the Convex and file transfer between the two systems was likewise under ACCESS

control. Over on the Convex, the FORTRAN77 code was studied, re-written, and recompiled using the usefully verbose vector compiler and in-house utilities available on that machine. Movie software by Rick Ottolini was the incomparable means for studying results in a dynamic and graphical way on a Rastertek monitor. The electronic network available through the Convex is an invaluable means of communicating with remote colleagues.

ACKNOWLEDGMENTS

Dan Rothman continues to be a source of lively electronic comment and criticism, and it would be difficult to imagine this paper coming together in this report without the help of that critical mass, my colleagues in the Stanford Exploration Project.

FURTHER READING

While there is yet no standard text on lattice gases, Wolfram (1986a) provides an excellent introduction, an important critical bibliography, and a useful collection of recent papers leavened with otherwise unpublished material. The same author has also recently published the first of several papers designed to review the field of fluid automata (Wolfram, 1986b).

REFERENCES

Frisch, U., Hasslacher, B., and Pomeau, Y., 1986, Lattice gas automata for the Navier-Stokes equation: Phys. Rev. Lett. 56 1505

Hardy, J., de Pazzis, O., and Pomeau, Y., 1976, Molecular dynamics of a classical lattice gas: Phys. Rev. A 13 1949

d'Humieres, D., Lallemand, P. and Frisch, U., 1986, Lattice gas models for 3D hydrodynamics: Europhys. Lett. 2 (4) 291-297

Muir, F., 1986, An automaton for propagating waves through heterogeneous media: SEP–50

Wolfram, S., 1986a, Theory and applications of cellular automata: World Scientific

Wolfram, S., 1986b, Cellular automaton fluids 1: basic theory: J. Stat. Phys. 45 3/4 471ff