

els are in general equivalent to single-species models operating on separate lattices. Colored collision rules couple the lattices so that information can be transferred between them at different time scales. Certain statistical-mechanical phenomena such as phase transitions can be done this way.

By altering the rule domain and adding gas species with distinct speeds, it is possible to add independent energy conservation. This allows one to tune gas models to different equations of state. Again, we gain no fundamental insight into the development of large collective models by doing so, but it is useful for applications.

In using these lattice gas variations to construct models of complex phenomena, we can proceed in two directions. The first direction is to study whether or not complex systems with several types of coupled dynamics are described by skeletal gases. Can complex chemical reactions in fluids and gases, for example, be simulated by adding collision rules operating on colored multi-speed lattice gases? Complex chemistry is set up in the gas in outline form, as a gross scheme of closed sets of interaction rules. The same idea might be used for plasmas. From a theoretical viewpoint one wants to study how much of the known dynamics of such systems is reproduced by a skeletal gas; consequently both qualitative and quantitative results are important.

Exploring Fundamental Questions. Models of complex gas or fluid systems, like other lattice gas descriptions, may either be a minimalist description of microphysics or simply have no relation to microphysics other than a mechanism for carrying known conservation laws and reactions. We can always consider such gas models to be pure computers, where we fit the wiring, or architecture, to the problem, in the same fashion that ordinary discretization schemes have no relation to the microphysics of the problem. However for lattice gas models, or cellular-

continued on page 214

REYNOLDS NUMBER *and* Lattice Gas Calculations

The only model-dependent coupling constant in the Navier-Stokes equation is the viscosity. Its main role in lattice gas computations is its influence on the Reynolds number, an important scaling concept for flows. Given a system with a fixed intrinsic global length scale, such as the size of a pipe or box, and given a flow, then the Reynolds number can be thought of as the ratio of a typical macrodynamic time scale to a time scale set by elementary molecular processes in the kinetic model.

Reynolds numbers characterize the behavior of flows in general, irrespective of whether the system is a fluid or a gas. At high enough Reynolds numbers turbulence begins, and turbulence quickly loses all memory of molecular structure, becoming universal across liq-

uids and gases. For this reason and because many interesting physical and mathematical phenomena happen in turbulent regimes, it is important to be able to reach these Reynolds numbers in realistic simulations without incurring a large amount of computational work or storage.

Some simple arguments based on dimensional analysis and phenomenological theories of turbulence indicate, at first glance, that any cellular automaton model has a high cost in computer resources when simulating high-Reynolds-number flows. These arguments appeared in the first paper on the subject (Frisch, Hasslacher, and Pomeau 1987) and were later elaborated on by other authors. We will go through the derivation of some of the more severe constraints on simulating high-Reynolds-number flows with

cellular automata, then discuss some possible ways out, and finally estimate the seriousness of the situation for a realistic large-scale simulation.

The turbulent regime has many length scales, bounded above by the length of the simulation box and below by the scale at which turbulent dynamics degenerates into pure dissipation, the so-called dissipation scale. We focus on these extreme scales and, with a few definitions, derive a bound on the computational storage and work needed for simulating high-Reynolds-number flows with cellular automata.

The Reynolds number R is usually defined not in terms of time but simply as $R = vL/\nu$ where L is a characteristic length, v is characteristic speed, and ν is the kinematic shear viscosity. One sees immediately why calculating viscosity functions for particular models is important. It is the only variable one can adjust in a flow problem, given a fixed flow in a fixed geometry. First, we calculate a rough upper bound on Reynolds numbers attainable with lattice models. If the speed of sound in the lattice gas is c_s and the spacing between lattice nodes is ℓ , then by definition the kinematic viscosity $\nu \geq c_s \ell$. Now viscosity estimated this way must agree with that fixed by the scale of hydrodynamic modes. Given a global length L and a global velocity V associated with these modes, $R = VL/\nu$ at best. In terms of the Mach number ($M = V/c_s$), the Reynolds number is equal to ML/ℓ . But M also characterizes fluid flow, and L and ℓ are model-dependent. In a lattice gas we can relate the ratio L/ℓ to the number of nodes in the gas simulator, namely $n = (L/\ell)^d$, where d is the space dimension of the model. Therefore, the number of nodes in a lattice model must grow at least as $n \sim (R/M)^d$. Computational work is the number of lattice nodes per time step multiplied by the number of time steps required to resolve hydrodynamical fea-

tures. This is $L/\ell M$ steps (to cross the hydrodynamical feature at the given Mach number), and so we find the computational work is of order R^{d+1}/M^{d+2} . For a so-called normal simulation based on the usual ways of discretizing the Navier-Stokes equation, the growth in storage is roughly proportional to one power lower in the Reynolds number than the growth in storage for the lattice gas. So at first it seems that simulating high-Reynolds-number flows by lattice gas techniques is costly compared to ordinary methods.

This argument is not only approximate; it is also tricky and must be applied with great care. The normal way of simulating flows escapes power-law penalties by cutting off degrees of freedom at the turbulence-dissipation scale, which the lattice gas does not do. The gas computes within these scales and so wastes computational resources for some problems. Actually computation of these very small scales is the source of the noisy character of the gas and is responsible for its power to avoid spurious mathematical singularities. One way around this is to find an effective gas with new collision rules for which the dissipation length scales are averaged out. A possible technique uses the renormalization group, but it is useful only if the effective gas is not too complex and has the attributes that made the original gas attractive, including locality. Work is going on at present to explore this possibility, and it seems likely that some such method will be developed.

The more serious consideration is what happens in a realistic large-scale simulation, and here we will find the lattice gas does very well indeed.

First, we note that a dissipation length l_d with the behavior $l_d \rightarrow \infty$ as $R \rightarrow \infty$ is actually required to guarantee the scale separation between the lattice spacing and the hydrodynamic modes that is necessary to develop hydrodynamic behavior.

The actual Reynolds number in lat-

tice gas models is much more complex than in normal fluid models. An accurate form is $R = Lv g(\rho_0)/\nu(\rho_0)$, where v is an averaged velocity and the fundamental unit of distance (the lattice spacing ℓ) and the fundamental unit of time (the speed required to traverse the lattice spacing ℓ) have been set to 1. To remain nearly incompressible, the velocities in the model should remain small compared to the speed of sound c_s , but c_s in lattice gases is model-dependent. So we factor the Reynolds number into model-dependent and invariant factors this way: we define $\hat{R}(\rho_0) = c_s (g(\rho_0)/\nu(\rho_0))$ so that $R = M L \hat{R}(\rho_0)$. The value of \hat{R} depends critically on the model used. In two dimensions it ranges from 0.39 to about 6 times that, depending on the amount of the state table we want to include. For the three-dimensional projection of the four-dimensional model, it is known that \hat{R} is about 9.

By repeating essentially the same dimensional arguments, only more carefully, we find that the dissipation length $l_d = (M \hat{R})^{-1} R^{-1/2}$ for two dimensions and $l_d = (M \hat{R})^{-1} R^{-1/4}$ for three dimensions.

For a typical simulation in three dimensions, we take $M = 0.3$ for incompressibility, $\hat{R} \sim 9$, and $L = 10^3$, which is a large simulation, possible only on the largest Cray-class machines. Then l_d is about three lattice spacings, and the simulation wastes very little computational power. The subtle point is that the highly model-dependent factor \hat{R} is not of order 1, as is usually estimated. It depends critically on the complexity of the collision set, going up a factor of 20 from the elementary hexagonal model in two dimensions to the projected four-dimensional case with an optimal collision table.

There is a great deal of work to be done on the high-Reynolds-number problem, but it is clear that the situation is complicated and rich in possibilities for evading simple dimensional arguments. ■