PHYSICAL REVIEW A VOLUME 36, NUMBER 7 OCTOBER 1, 1987

Microscale hydrodynamics: Discrete-particle simulation of evolving flow patterns

D. C. Rapaport
Department of Physics, Bar-Ilan University, 52 100 Ramat-Gan, Israel
(Received 2 March 1987)

The technique of molecular-dynamics simulation—in which the equations of motion of a system of interacting particles are solved numerically to yield the temporal evolution of the system—is used in a study of the flow of a two-dimensional fluid past a circular obstacle. The flow is observed to develop with time, passing through a series of well-defined patterns that bear a striking similarity with flow patterns observed experimentally in liquid and gas flow; the patterns include stationary eddies, periodic shedding of vortices, and a vortex street characterized by a Strouhal number close to the experimental value. Very large systems—by current molecular-dynamics standards—need to be used in order to accommodate the obstacle and the region occupied by the structured wake, and the present work includes the largest such simulations carried out to date. Though more extensive work is called for, the results suggest that continuum hydrodynamics is applicable down to much shorter length scales than hitherto believed, and that the molecular-dynamics approach can thus be used to study certain kinds of hydrodynamic instabilities.

I. INTRODUCTION

Low-Reynolds-number hydrodynamics embraces a class of fluid dynamical problems that are responsible for a fascinating kaleidoscope of flow patterns. The different flows, both static and time varying, are obviously all solutions to the equations of continuum hydrodynamics, but there has been little by way of analytical progress towards predicting the conditions under which each kind of flow might be expected. The definition of the Reynolds number for fluid flowing past an obstacle is $\mathcal{R}=UL/\nu$, where U is the stream velocity, L the characteristic linear dimension of the obstacle, and ν the kinematic viscosity. If the fluid is water, and $L=10$ cm, a value $\mathcal{R}=1000$ is attained when U is only 1 cm/s, while for air the corresponding U is 15 cm/s.[1] Thus it is apparent that flow studies at low $\mathcal{R}$, in particular for $\mathcal{R}<100$, are not typical of flows normally encountered in engineering applications. However, in order to understand the stability or otherwise of the various possible solutions to the hydrodynamic equations, and to learn how periodic solutions eventually lead to chaotic behavior—in this context, turbulence—it is necessary to begin at the low end of the Reynolds-number scale. Therein lies much of the motivation for low-$\mathcal{R}$ hydrodynamics.

In this paper we describe a new approach to the problem of low-$\mathcal{R}$ flow in which the fluid is represented by means of its primary constituents—a collection of mutually interacting particles. The flow is determined by a direct numerical solution of the Newtonian equations of motion—the well-known technique of molecular-dynamics simulation—and subsequent course graining to permit visualization of the bulk motion. A preliminary account of this work has appeared elsewhere.[2]

One might reasonably ask the question why such an approach was not pursued in the past. The answer is twofold: First, there was probably little reason to believe that inhomogeneous time-dependent flows limited to the microscopic length and time scales achievable by molecular dynamics would resemble those observed at the macroscopic level; second, the size of the minimal microscopic system capable of exhibiting unambiguous, well-defined flow patterns turns out to be some 2 orders of magnitude larger than systems typically employed in molecular-dynamics calculations, and studies of this nature have therefore had to wait until computers of adequate power became readily available.

A brief background to the subject is presented in Sec. II; both experimental studies of flow past an obstacle—the most common being a long circular cylinder—and various numerical approaches are briefly surveyed. Section III is devoted to methodology; in addition to outlining the extensions to standard molecular dynamics required to deal with the flow problem, a discussion is included of how the problem was decomposed to run in parallel on a network of coupled processors. The methods used to aid flow visualization are discussed in Sec. IV. The resulting flows, which look remarkably like those seen experimentally, are described and illustrated in Sec. V. Section VI attempts a critical review of the molecular-dynamics approach and its relationship to "real" hydrodynamics. Finally, Sec. VII outlines the possible future of molecular dynamics in the field of fluid dynamics with suggestions as to how the present work might be furthered.

Simulation of complex flow at the discrete particle level, should it prove to be a viable approach, may serve as a means of accessing the subtle correlations that must surely be present in the seemingly chaotic microscopic motion, and which are responsible for the intricate, well-structured, and beautiful flow patterns actually observed. Experimentally it is difficult, if not impossible, to probe the details at this level, and it is here that the

molecular-dynamics approach could serve as a useful adjunct to the study of real flows.

II. BACKGROUND

A. Experimental observation

Flow instabilities are readily observed at low $\mathcal{R}$ in fluids streaming past bluff obstacles; whereas much of hydrodynamics is concerned with streamlining in order to minimize the disturbance to flow, bluff bodies are used with exactly the opposite intent, and their effect on flow has been extensively studied.[3] Obstacle shapes include thin plates normal or inclined to the stream, cylinders of circular and elliptical section, and spheres. The obstacle that will be emphasized here is a circular cylinder that is sufficiently long (the axis is normal to the flow) to permit the problem to be regarded as two dimensional. There is, however, a great deal in common among the flows due to differently shaped obstacles. As a matter of convenience, and without affecting the physics, it is often the obstacle rather than the fluid that is set in motion, usually impulsively from rest.

The observed flows at increasing Reynolds numbers[1,4] can be summarized as follows: At very low $\mathcal{R}$ the flow is Stokes like, with the component of the velocity field parallel to the stream always in the forward direction. The initial departure from forward flow occurs at $\mathcal{R}\approx 5$ at which point two oppositely rotating eddies start to develop adjacent to the downstream face of the cylinder. As a consequence, the flow direction behind the center of the cylinder is reversed. The downstream length of the eddies increases with $\mathcal{R}$, reaching approximately twice the cylinder diameter at $\mathcal{R}\approx 40$. Oscillations in the wake downstream are seen, beginning close to $\mathcal{R}=34$. When $\mathcal{R}$ reaches a value in the range 55–70 (the value cannot be established more precisely than this), the eddies begin to oscillate laterally, with a rotating region of fluid being shed each half-period. These regions are evenly spaced in two parallel rows and form what is known as the von Kármán vortex street. The pattern propagates downstream with a velocity below that of the unaffected flow and tends to spread gradually. The shedding frequency increases with $\mathcal{R}$, and above $\mathcal{R}\approx 100$ the eddy pair appears to vanish. Further subtle changes occur as $\mathcal{R}$ is increased beyond this value, including the appearance of additional frequencies, eventually leading to turbulence, but this regime is already beyond the scope of the present study

The flow details are best conveyed pictorially, and photographs of eddies and vortex streets are widely available.[1,5] The photographs clearly show the variety of flow patterns that occur, but due to the relatively coarse resolution of the visualization techniques used (dye, bubbles, injected particles), not all the subtle details of the flow are revealed; for example, the boundaries (separatrixes) between regions of fluid undergoing different forms of motion.[6] The existence of complex flow patterns has been known for a long time: Eddy formation in the lee of an obstruction is included among the sketches of Leonardo da Vinci.[7] The phenomenon of vortex shedding is responsible for the familiar Aeolian sounds[8] heard, for example, when wind blows past an overhead cable.

B. Theoretical approaches

The nonlinear hydrodynamic equations do not admit an analytic treatment at the Reynolds numbers of interest and it is therefore necessary to resort to numerical solution. There have been inexplicably few attempts to produce solutions that exhibit time-dependent behavior. In the two computations that have looked at vortex shedding in two-dimensional flow past a circular obstacle,[9,10] it proved necessary to artificially perturb the system to break up the stationary eddy pair that would otherwise have been the preferred solution. Without this injection of vorticity to produce time-varying flow the eddies persist to unreasonably high $\mathcal{R}$ values, at least as far as $\mathcal{R}=1000$.[11] This serves as an excellent example of the need not only to obtain a numerically reliable solution to the equations, but to ensure that the solution is the physically appropriate one, and not one that has been continued beyond its region of natural stability.[12]

The computations just cited employed nonslip boundaries. This is not an essential ingredient, however; a recent calculation for flow past an ellipsoid showed that eddies can develop even with slip boundaries.[13] It is therefore plausible to conclude that the appearance of eddies is due to vorticity generated by the mere presence of a bluff obstacle, irrespective of the nature of the boundary conditions. This statement is not true, of course, for streamlined bodies where the nonslip requirement is of unquestionable importance. One rather puzzling aspect of this particular computation is that the eddies eventually begin to decrease in size as $\mathcal{R}$ grows; it is tempting to speculate that this may be another instance of a solution that is no longer appropriate, and that perturbing the flow will again lead to vortex shedding.

Alternative computational approaches abandon the continuous density and velocity fields in favor of assemblies of discrete particles that are treated at varying levels of realism. The molecular-dynamics (MD) technique—to be discussed in Sec. III—aims at the most precise representation of the motion of the fluid molecules. Due to this surfeit of detail the computational demands of this approach are heavy. One means of reducing detail is to sample only some of the collisions the molecules experience;[14] this technique is more appropriate for dilute gases than dense fluids, fails to conserve angular momentum and, consequently, produces incorrect results for obstructed flow.[15] Another recently introduced technique is based on cellular automata[16,17] and has succeeded in producing flow patterns such as the vortex street.[18,19]

III. MOLECULAR DYNAMICS METHODOLOGY

A. Basic technique

The molecular-dynamics technique has been used in a wide variety of equilibrium and nonequilibrium contexts over the past thirty years.[20–23] While a number of specialized techniques have been developed to improve the computational effectiveness of the technique, MD actual-

ly amounts to no more than the solution of the coupled equations of motion for the particles that make up the system. The result of an MD computation is a set of particle trajectories; appropriate averaging leads to quantities corresponding to the thermodynamic and transport properties of the system.

The wide variety of fluids that are Newtonian in nature and which, when compressibility effects are minimized, exhibit flows that depend only on $\mathcal{R}$ (for given boundary conditions)—typical extremes being air and water—suggests that the details of the intermolecular force are irrelevant for an initial MD approach. This implies that the principal criterion for the choice of interaction potential should be one of computational efficiency; a short-range repulsive interaction between isotropic particles is the ideal choice. Whether this is realized by means of a step potential so that the particles are hard elastic disks in two dimensions (spheres in three), or as a smooth function that diverges as the particles approach too closely and tapers off to zero at some finite particle separation, is of little consequence to the results, but can influence the efficiency of the MD implementation.

The algorithms used for continuous and step potential are very different; the former is simply a straightforward numerical integration of the second-order equations of motion with a constant time step, while the latter is based on a scheme for scheduling particle collisions with the system advancing in time from one collision to the next.[23,24] The relative efficiency of the algorithms is strongly dependent on the architecture of the computer; on a scalar processor the step potential can be more effective than the continuous, but the algorithm is not suited for implementation on either vector or parallel processors. Thus the choice of processor (see below) dictated the use of a continuous potential. The potential function has the form

$$V(r)=\begin{cases} r^{-12}-r^{-6}+\frac{1}{4}, & r<r_c=2^{1/6} \\ 0, & r\geq r_c \end{cases}.$$

where r is the interparticle separation. This is a shifted and truncated version of the Lennard-Jones interaction in which the attractive tail has been eliminated leaving only a repulsive core. A third-order predictor-corrector method[25] is used to solve the coupled equations of motion.

Although the equations of motion are reversible, the numerical methods and the finite precision arithmetic of the computer introduce noise into the solution, so that for practical purposes the system is irreversible. One manifestation of this effect is a lack of strict energy conservation, and for closed systems a periodic rescaling of the particle velocities serves to eliminate the small but persistent energy drift. The fluid flow simulation corresponds to an open system so that the issue of energy conservation does not arise; individual trajectories are still subject to error, but the fact that realistic flow patterns emerge from the simulation (Sec. V) is an indication that this does not present a serious problem. This in turn prompts the question of how precise the dynamical computations need to be—how much error can be tolerated in recording the coordinates and velocities, or in computing the accelerations at each time step? Reduced precision could lead to faster computation, but at some stage the results must cease to have any resemblance to the correct solution. Exploration along these lines is warranted.

B. Introducing the obstacle

In continuum fluid dynamics the distinction is made between slip and nonslip boundaries. The origin of the nonslip condition lies primarily in the roughness of the surface; since the length scales that characterize this irregularity are significant, possibly exceeding even the entire size of the simulated system, it is not apparent that roughness can be represented directly, except possibly in stylized form using, for example, fractals.[26] There exist alternative possibilities for achieving the effects of roughness while actually using a smooth boundary. One such technique is to randomize the velocity of a particle at its closest approach to the boundary while keeping the kinetic energy unchanged. An even less drastic technique is to just randomize the sign of the tangential velocity component. Both approaches produce the desired effect. A scheme to be avoided in this kind of study is the randomization of velocity that also adjusts the kinetic energy to correspond to a prescribed boundary temperature; this would introduce spurious effects since the obstacle then acts as a thermal source or sink.

The other choice is to adopt slip boundaries from which the colliding particles rebound elastically. In view of the discussion earlier, this is not unreasonable from the hydrodynamic point of view. Freedom from the nonslip requirement permits the introduction of an obstacle that can be expressed in terms of a potential function and incorporated directly into the equations of motion. In the interest of simplicity the potential used to represent the interaction between particle and obstacle is the same function $V(r)$ used for the interparticle interaction, but the distance r is now taken to be between the particle and the closest point on the obstacle boundary. For a circular obstacle of diameter D, particles within a distance $D/2+r_c$ of the center experience a repulsive force; the effective diameter of the obstacle is increased slightly, but since $D \gg r_c$ the change is negligible. Circular and rectangular obstacles are easily handled, and more complex shapes can also be used at the cost of a little extra computation.

C. Closing the open system

The system is an open one, in the sense that to maintain flow past a fixed obstacle, particles must be injected upstream at an appropriate rate and removed as they exit downstream. In order to conserve particle number, departing particles are reintroduced immediately. Provided there are no significant density inhomogeneities close to the downstream limit of the system, periodic boundaries can be used; however, flow patterns in the wake of the object must be eliminated before the particles reenter the system, a result achieved by the simple

expedient of injecting each particle with a constant velocity in the stream direction on which is superimposed a random thermal contribution corresponding to the ambient temperature. To maintain pressure, particles at the upstream and downstream extremities of the system are allowed to interact, as is normally the case with periodic boundary conditions; there is no indication that the downstream flow influences the reentering fluid in any significant fashion. The boundaries parallel to the flow are also periodic, but require no special attention.

Ideally, the particles entering and leaving the system should be completely isolated from each other. This could be achieved by introducing suitably shaped potential barriers at each end of the system whose task is to sweep out departing particles and force reentering particles back into the system under pressure. Such an approach would keep the particles confined to the allowed region and at the same time maintain a gap between the particles at either end. There is also the need for an external "gravitational" field to maintain the flow. Downstream flow momentum is dissipated in the collisions with the obstacle, and in the absence of a field the initial flow would quickly cease. The field strength must be determined by trial and error; too weak a field allows the flow to eventually die out, while a too strong field leads to excessive velocities. If a moving rather than a stationary obstacle had been chosen, the role of the field would have been to prevent the fluid being dragged along with the obstacle. The pressure head created by the potential barriers mentioned above could also substitute for the gravitational field.

D. Enhancements

Two enhancements were introduced into the MD algorithm in order to accelerate the computation. The first is the subdivision of the system into cells in order to convert the $O(N^2)$ calculation required to locate interacting pairs of particles into one requiring only $O(N)$ effort, where N is the number of particles in the system.[24] Particles are assigned to this two-dimensional cell array on the basis of their coordinates; if the cell edge exceeds r_c in length, then interacting particles can only lie in the same or immediately adjacent cells. The optimal cell size has to be selected empirically, but the typical mean occupancy should be of order unity. Use of cells also aids in locating particles that are within interaction range of the obstacle. The second enhancement is the elimination of the interaction calculation in favor of a table lookup based on the square of the interparticle distance.

No attempt was made to explore the use of various synthetic dynamical approaches that have been used recently in other nonequilibrium contexts.[27] It is not clear how the use of constant temperature or pressure ensembles would improve the results; given the difficulty of measuring thermodynamic properties in the highly inhomogeneous flows encountered in this study, as well as the spurious physical effects that can arise if the synthetic approach is not used with extreme caution,[28] it is probably safest to use straightforward Newtonian dynamics at this juncture.

E. Parallel processing for large problems

It has long been acknowledged that one way of achieving increased computational power is to resort to parallelism—where $K(\geq 2)$ processors cooperate in performing a calculation in a fraction of the time required by a single processor; in the optimal case the time is proportional to $1/K$.[29] Many computations cannot achieve this idealized goal due to the communication overheads required in sharing data between processors and because some algorithms contain steps that cannot be executed in parallel; however, an MD simulation that involves only short-range interactions turns out to be an excellent example of a problem that can be distributed over a set of communicating computers with near-optimal performance. The modifications to the basic MD approach needed to achieve this goal are described in the remainder of this section.

Given an array of processors capable of communicating with one another, either directly or indirectly via intermediary processors, the problem is one of evenly distributing the computational load while ensuring that only a minimum of data need be exchanged among processors. There are two distinct ways of dividing up the problem; either by making each processor responsible for a particular subset of particles, or by giving each processor the task of handling the computing for a separate spatial subregion of the system. The former scheme[30] is more appropriate for problems involving long-range interactions since a complete force calculation requires that all particle coordinates be made available to all processors, implying a considerable amount of data movement between processors. The latter scheme is ideal for short-range forces since the only coordinate information that must be communicated is associated with particles close to the subregion boundaries which lie within interaction range of particles in the adjacent subregions; for short-range forces the numbers of particles involved are small and thus data transfer is kept to a minimum. This approach was adopted for the present study.

The two-dimensional system can be subdivided into K equal slabs of width s so that particles whose x coordinate lies in the range $(k-1)s \leq x < ks$ are assigned to processor k $(1 \leq k \leq K)$. The alternative is to use a two-dimensional subdivision, but since extra computation and data transfer is required it is more suited to larger K values than used here. The one-dimensional subdivision into slabs, together with the imposed periodic boundaries, implies that the optimal processor network has the connectivity of a ring.

The stages of the predictor-corrector algorithm for a single time step, modified for parallel processing, are as follows.

(1) Do predictor step; all K processors in parallel.

(2) Transfer particles that have exited slab: All information about those particles i in processor k for which $x_i \geq ks$ is transferred to processor $k+1$, and those with $x_i < (k-1)s$ to processor $k-1$; the periodic condition implies that slabs 1 and K are adjacent.

(3) Copy coordinates of particles close to slab boundary: Coordinate data for particles in processor k for which $x_i \geq ks - r_c$ are duplicated in processor $k+1$, and for $x_i < (k-1)s + r_c$ in processor $k-1$; this step allows the interaction computations to proceed independently.

(4) Do interaction calculations; all K processors in parallel.

(5) Do corrector step; all K processors in parallel.

The transfer and copy operations proceed sequentially around the ring, but the work necessary to organize the data being communicated can also be carried out in parallel.

The largest simulations of the present study used four Floating Point Systems FPS-264 scientific processors communicating via a common shared memory.[31] The processors were under the control of an IBM 3081 mainframe which played only a supervisory role and was not involved in the computations. The typical computation rate with 4×10^4 particles per processor (using a mixture of Fortran and assembly language) was 1700 time steps per hour, irrespective of the overall system size. The communication overhead never amounted to more than about 2%. An indication of the critical nature of the communications can be gathered from the fact that when data transfers proceeded via the mainframe instead of the shared memory, the rate dropped to 1100 steps per hour for four processors.

IV. FLOW VISUALIZATION

The nature of the output of these simulations defies concise summary and really calls for an album of images as is the case experimentally. An even better medium is videotape as this performs the temporal interpolation necessary to follow developing patterns. Since this is not feasible, the results are presented as a series of images of flow patterns accompanied by graphs showing quantitative aspects of the flow.

The simulations have the capability of generating vast amounts of detailed data, and a certain amount of filtering is necessary at the outset to avoid being overwhelmed. This was achieved by dividing the region into a 60×60 cell grid and computing the mean velocity and number of particles in each cell (note that this has no bearing on the dynamics). The evaluation was carried out every fifth time step; the average of 200 such instantaneous images represents a single block of filtered data, amounting essentially to a "time exposure." Further combining of these blocks of data was used to produce the images shown, but no additional data smoothing was applied and the images represent the raw data as actually obtained.

The optimal use of static images to convey a description of complex flow would involve trajectory plots of suitably injected test particles. Unfortunately, the temporal resolution of the filtered data is too coarse to support this approach. An almost as effective an alternative is to plot flow fields, either magnitude and direction, or just direction; the reader's eye can make the necessary inferences. Both are used here.

One kind of plot consists of arrows conveying information about flow velocity, direction, and particle density in each grid cell. Arrow length is related to velocity magnitude, but because flow rate in the lee of the obstacle is a great deal lower than average, and simple proportionality would lead to very short arrows in this area, the arrow lengths are prescribed by the relation *length* $\propto 1 + v/v_{\max}$, where $v_{\max}$ is the largest velocity value in the plot. Arrow head size is proportional to local density. Deviations from downstream flow by more than 10° are highlighted by thicker arrows.

The other kind of display conveys information about flow direction only and is a rough approximation of a streamline plot. As used here, only the central 20×40 cell area just downstream of the obstacle center is shown; this is where the interesting behavior arises. The continuity of the flow lines is improved, especially in regions of high curvature, by dividing the line segment in each cell into three parts with each outer part having a direction that is a weighted average of the flow direction within the cell itself and that of the cell to which the line points. The scheme is arbitrary but produces the desired effect. Flow deviations are again highlighted.

V. RESULTS

A. Summary of runs

The large-scale simulations described here represent the culmination of a series of runs with systems of monotonically increasing size in an effort to observe well-defined, and possibly even time-varying, flow patterns. With smaller systems—typically containing just half as many particles—only a tantalizing hint of what is happening can be seen, since there is insufficient space around the circular obstacle (which has a minimal size not significantly less than used here) for the flow pattern to develop owing to boundary interference. Even the largest system is not fully exempt from this limitation. Other obstacle shapes were also explored prior to attempting the circle—in particular, linear barriers of various thickness perpendicular to the stream, with straight or semicircular ends. For a sufficiently high ratio of length to thickness interesting flow patterns could be observed with as few as 10^4 particles. However, it is the circular obstacle that presents the greatest challenge, since all available evidence indicates that the flow instabilities only begin at some nonzero critical Reynolds number $\mathcal{R}_c$, the estimated value being close to 5. An object with a boundary whose direction does not vary continuously (i.e., with corners) forces flow separation to occur at the points of discontinuity since flow cannot abruptly change direction; the value of $\mathcal{R}_c$ in this case—which includes the thin plate—is essentially zero.[1] For an obstacle without corners the point of flow separation depends on $\mathcal{R}$, provided of course that $\mathcal{R} > \mathcal{R}_c$. A systematic study of obstacle shape and size is a possible extension of the present work.

Dimensionless units are used in the MD simulations. The actual values traditionally used correspond to the length and energy typical of liquid argon,[32] although a

TABLE I. Summary of runs. The units employed are defined in the text.

Run number	Particles (10^3)	Field (10^{-4})	Initial flow velocity	Time steps (10^3)	Final mean flow velocity
1	159	0.17	0.1	25	0.09
2	168	0.66	0.2	66	0.23
3	168	1.5	0.3	100	0.38
4	159	2.0	0.3	120	0.42
5	168	2.6	0.4	107	0.52

different choice of liquid would have little effect; thus the unit of length is 3·4 Å and that of time 0.31 ps. The systems used in the full-scale runs contained $1.6-1.7\times10^5$ particles and the runs were continued for as many as 1.2×10^5 time steps (1.1 ns); the size of the time step is 0.03 ($\approx10^{-14}$ s). The key run parameters appear in Table I. The bigger system is, to the author's knowledge, the largest MD simulation to have been carried out to date, a record that is unlikely to have a long lifetime. The mean density is 0.83 particles per unit area, a value substantially below freezing. Temperature, when it could be regulated—at the beginning of the run and when particles are reinserted—corresponds to a thermal velocity of 0.2. The circular obstacle has diameter 74, close to one-sixth the region width, and is positioned upstream of center. In real terms the region is practically a square of edge 1500 Å, so that the circle diameter is 250 Å.

The initial state is one of unidirectional flow, with random thermal motion superimposed. Under the influence of the field the flow velocity changes gradually and eventually stabilizes. Actual flow speeds are in the range 0.1–0.5 (MD units—the upper limit corresponding to approximately 550 m/s); while the speed of sound has yet to be determined for this system it should not be too different from the Lennard-Jones model of liquid argon, namely, 0.8.[32] Thus the Mach number is less than unity and shock waves should not present a problem.

B. Flow patterns

The patterns to be described below were observed in the course of run No. 4 (Table I). Other runs will be summarized subsequently. The flow plots with arrows are the averages over five successive data blocks (Sec. IV) corresponding to an elapsed interval of 150 MD time units (750 ps), while the plots showing direction only are based on two data blocks. Note that if the flow pattern is in the process of changing relatively rapidly, a certain amount of "blurring" will occur in the images. The images (Figs. 1 and 2) are to be followed in conjunction with the text.

At time $t=0$ the system begins with a uniform flow velocity plus thermal noise. By $t=60$ (the times correspond to the ends of the intervals) this has already changed into Stokes flow around the obstacle. The disturbed flow region grows until at $t=240$ there is a hint of imminent qualitative change. The first signs of reversed flow appear at $t=300$, and by $t=360$ a pair of

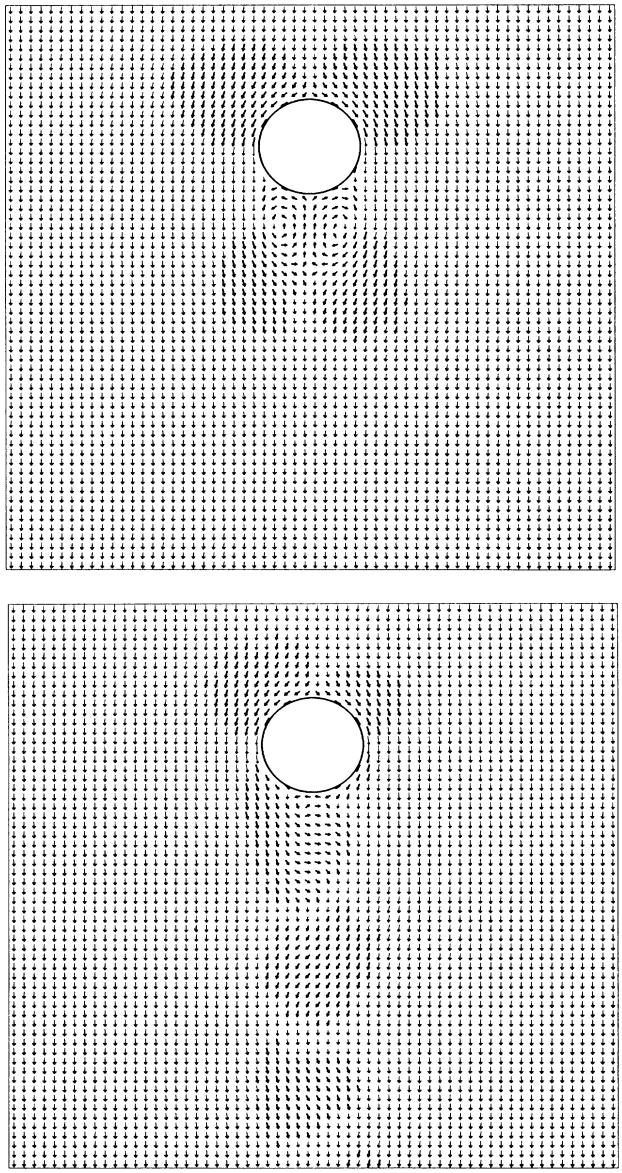

FIG. 1. Velocity plots showing the entire 60×60 cell flow field and obstacle position for run No. 4. The two frames correspond to times 600 and 3000 (MD units) and show, respectively, a well-formed eddy pair, and the fully developed wake oscillation with a very small eddy to the left of center. See text for interpretation of arrow size.

well-formed counterrotating eddies have established themselves just downstream of the obstacle. Their combined width is approximately the obstacle diameter D. The eddies stretch downstream, without any significant change of width, eventually reaching a length of approximately $1.4D$ at $t=840$. The reversed flow velocity amounts to only about 10% of the forward flow; just below the obstacle the density drops by about 25%, but upstream the increase is only 3% (see graphs in Fig. 4 below). The relative coarseness of the grid makes estimation of the separation angle of the forward-moving boundary layer difficult, but the value seems to be in the vicinity of 120°—as measured from the upstream stagnation point. The theoretical estimate for this angle under nonslip conditions is near 109°.[33]

A breakdown in the symmetric flow pattern is the

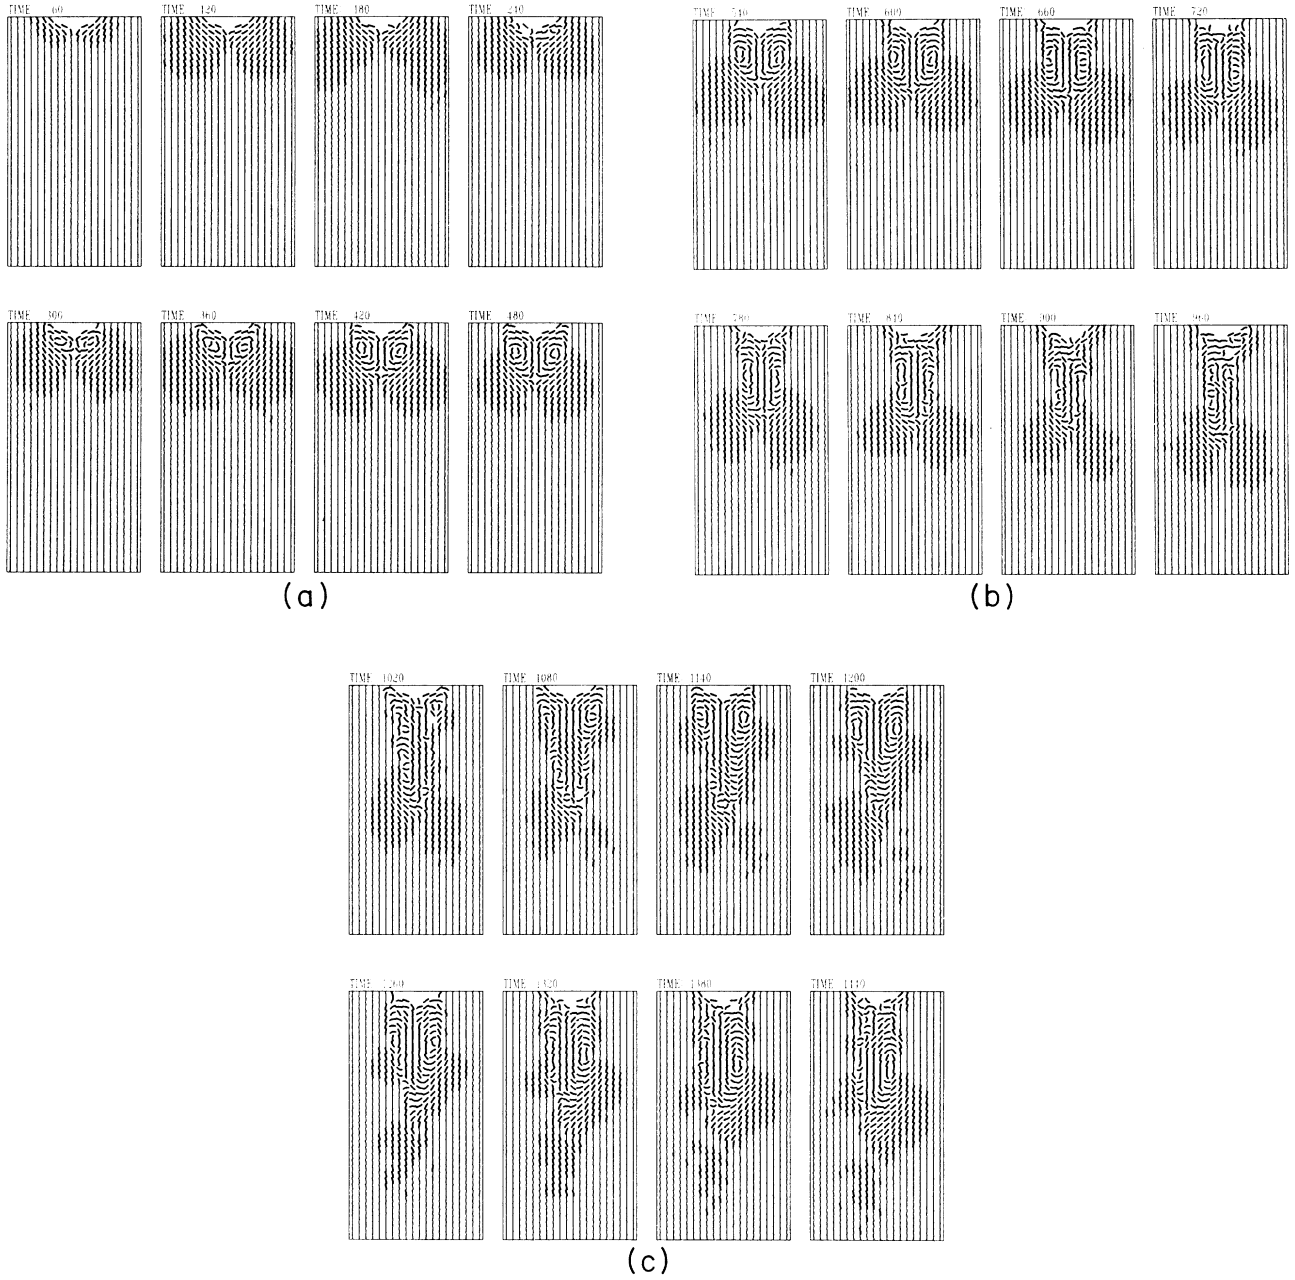

FIG. 2. Flow direction plots for run No. 4 covering a 20×40 cell region centered on the wake (there are no lines in the area covered by the obstacle). The first sequence of frames (a)–(c), for times 60 to 1440, shows the initial appearance and growth of the eddies, followed by symmetry breaking and restoration, and the beginnings of wake oscillation. The second sequence (d) and (e) covers times 2460 to 3360 and shows the traveling wave wake as well as the last remnants of the eddy.

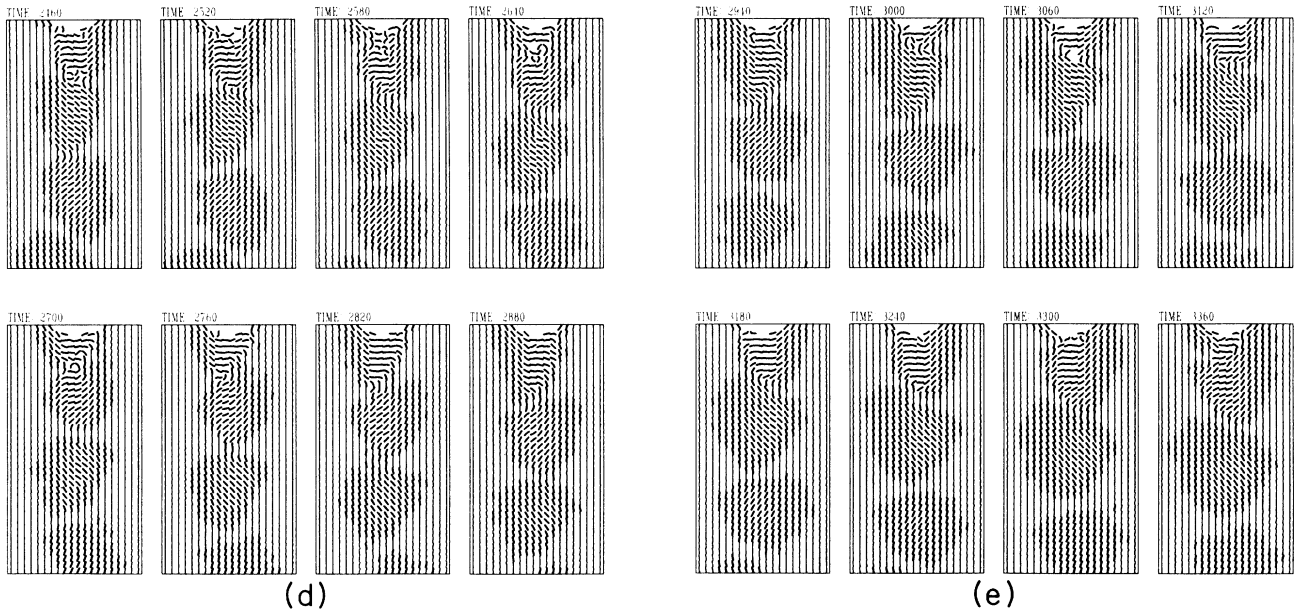

FIG. 2. (*Continued*).

next phenomenon to be observed. The left (counterclockwise-rotating) eddy begins to expand laterally at the expense of its companion which, by $t=960$, has all but disappeared, only to reappear shortly afterwards, with symmetry restored by $t=1200$. Just below the eddies, however, there is a region of leftward flow, a precursor to the transverse oscillations to appear in due course.

The subsequent flow behavior immediately downstream of the obstacle suggests that eddies grow alternately on either side of the equator (parallel to the stream) and are shed into the flow. All that remains of the eddy after it has traveled a short distance is a transverse contribution to the wake. The process is repeated several times until, by $t=2400$, there is an oscillatory wake stretching all the way to the downstream limit of the region, the remnant of previously shed eddies. The sinuous wake continues to propagate downstream although the eddies themselves diminish in size and completely disappear by $t=2800$; the wake oscillations then begin adjacent to the obstacle boundary. No new effects, other than a slight broadening of the wake, are noted before the run terminates at $t=3600$.

The sinusoidal-like wake does in fact correspond to a vortex street; it is merely a matter of selecting the correct frame of reference in which to view the flow, the appropriate frame being that in which the eye tracks the wake moving at a lower velocity than the undisturbed flow. Figure 3 shows the new appearance of the wake—two staggered rows of vortices moving downstream, with new vortices appearing each half-cycle.

C. Quantitative results

Quantitative aspects of the flow can also be deduced from the coarse-grained averages. Figure 4 shows local velocity, vorticity, and density deviations, at a series of evenly spaced transverse cross sections through the system at key stages in the flow development. The velocity shown is the modulus of the cell average, with the sign indicating forward or backward flow; vorticity is computed as the circulation integral approximated by a sum over the four grid cells surrounding the point; the densi-

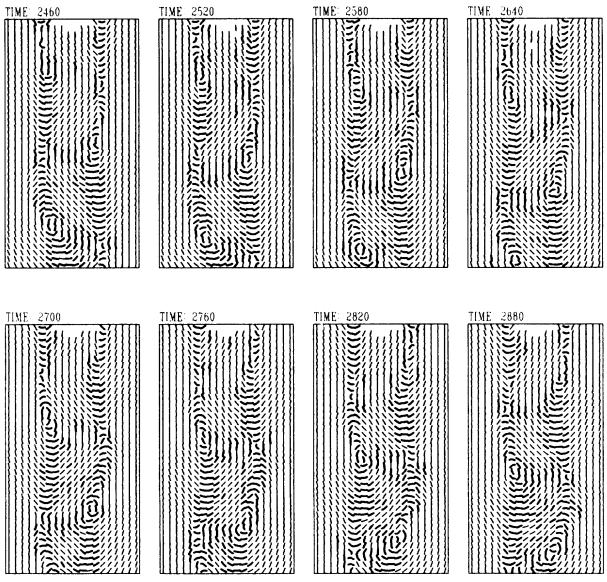

FIG. 3. Flow direction plots for times 2460 to 2880 already shown in Fig. 2 but from the perspective of an observer traveling downstream with the wake (the same spatial region is shown in each frame). The patterns correspond to a von Kármán vortex street.

ty deviation is simply the mean occupancy from which the overall average is removed.

The flow patterns strongly resemble what is observed experimentally on entirely different length and time scales. The resemblance would be even more significant, suggesting a common underlying mechanism, if various numerical characteristics of the flows also prove to be similar. Irrespective of the outcome of such a test, however, the very fact that a many-body simulation of this kind is at all capable of spontaneously generating coherent dynamic structures is a point worth noting.

Incompressible fluid flow past an obstacle is governed by the Reynolds number $\mathcal{R}=UL/\nu$, where L is set equal to D, the obstacle diameter. That the MD value of $\mathcal{R}$ should have anything in common with the experimental value is not obvious *a priori;* the substantial density variations, the large velocity gradients, and the extremely short length scales of MD all suggest this to be an unlikely occurrence. No attempt was made to estimate ν in the course of the computation (simulated viscometry), the value used being taken from a study of Couette flow;[34] this too creates problems since ν is temperature and density dependent and varies with shear rate.

Kinematic viscosity is defined as $\nu=\mu/\rho$ where μ is the shear viscosity and ρ density. The lowest shear rate of the Couette study is similar to that occurring here. Neglecting both the fact that shear rate is actually a function of distance from the obstacle (see Fig. 4), and the density variations, leads to the result $\nu=1.2$ (MD units). Taking $U=0.4$—the final flow speed—produces

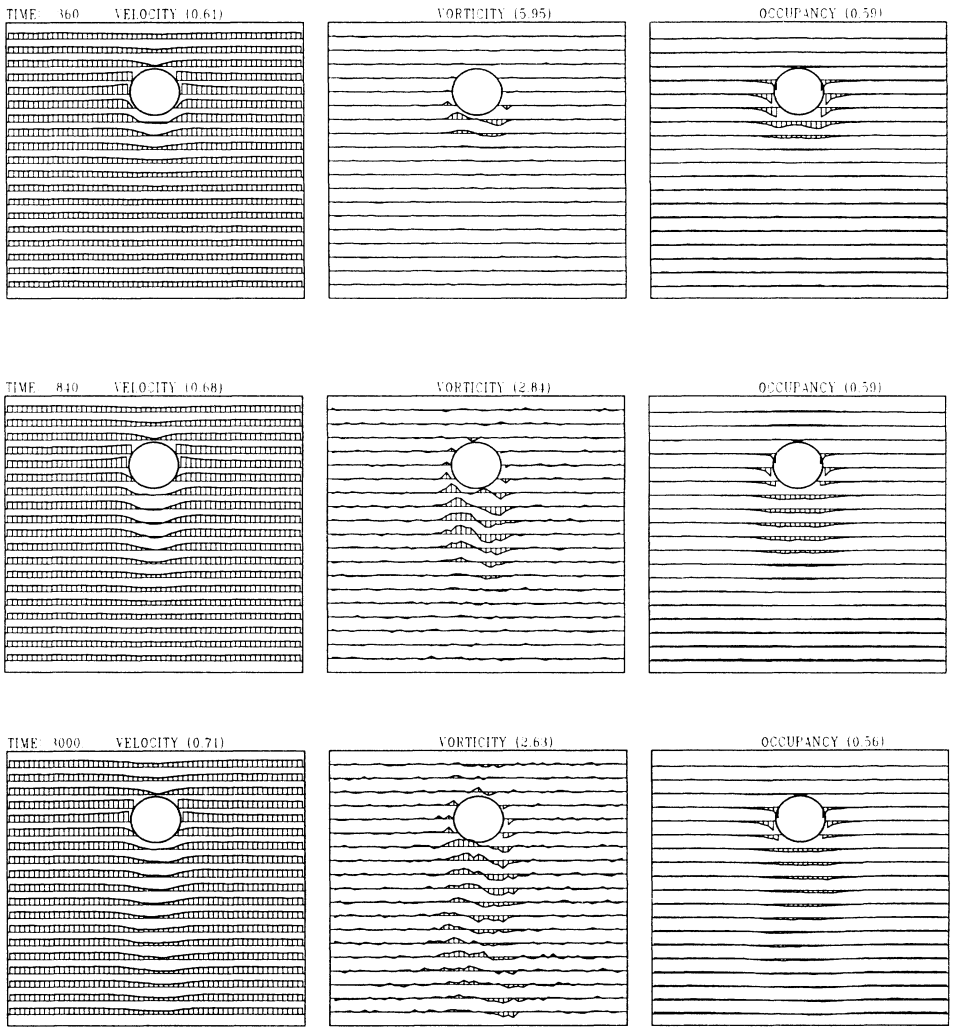

FIG. 4. Velocity, vorticity, and the deviation of cell occupancy from average (i.e., density variation) at a series of cross sections through the region. The graphs show the situation when the eddy pair has formed (time 360), when the eddy symmetry is about to break (840), and when the oscillatory wake has reached full length (3000). The values in parentheses are the maximum (minimum for occupancy) values shown in the particular plot. Note that the circle partially occludes cells immediately adjacent to it and exaggerates the occupancy drop in these cells.

the estimate $\mathcal{R}=25$, a value only a little below where wake oscillations first appear experimentally ($\mathcal{R}=34$). Given the above list of potential contributors to an erroneous $\mathcal{R}$ value, this rough estimate is satisfactory.

A quantitative measure of the temporal variation of the wake is provided by the dimensionless Strouhal number $\mathcal{S}=fD/U$, where f is the wake oscillation frequency. The Strouhal number can be used in two different ways,[1] one to specify an external perturbation, the other a measure of the spontaneous oscillation of the system. It is the latter usage that is intended here since the obstacle is maintained in a fixed position. The frequency is deduced from the wavelength and group velocity of the oscillatory wake; shortly after the appearance of oscillation these quantities have the values 220 and 0.24 (MD units) respectively, so that $f=1.1\times10^{-3}$ and thus $\mathcal{S}=0.2$. The corresponding low-$\mathcal{R}$ value measured experimentally is $\mathcal{S}=0.15$, a value which increases with $\mathcal{R}$ and asymptotes to approximately 0.2 for $\mathcal{R}>300$.[4] While the agreement is adequate and provides fairly convincing support that similar physical processes are operational at MD and hydrodynamic scales, it does raise the question that perhaps the value of $\mathcal{R}$ has been underestimated. Certainly the eventual disappearance of the eddies, leaving only the wake oscillation, is characteristic of flows with $\mathcal{R}>100$, but speculation along these lines is premature since it has yet to be demonstrated that a stable eddy pair (not the transient effect observed here) can be maintained at the length scales accessed by the present MD study.

D. Further runs

Other large-scale runs that were carried out are included in Table I. These assist in bounding the conditions under which flow structures appear but fall far short of being a systematic coverage of parameter space. In run No. 1 the velocity was down to about a quarter of the value in run No. 4; no deviation from Stokes (i.e., potential-like) flow could be seen. In run No. 2 the mean flow speed was just over half that of No. 4; a small eddy was noticed alternately forming and vanishing in the lee of the obstacle. Run No. 3 was only 10% slower and the behavior was very similar to that of No. 4, with the exception that the wake was slightly narrower. The mean flow rate in run No. 5 was 25% higher; the initial eddy pair was able to stretch right down to the end of the system before the breakup occurred, so the boundary clearly affected the behavior. Subsequent evolution was similar to No. 4, though with eddy oscillation still visible at the end of the run (Fig. 5).

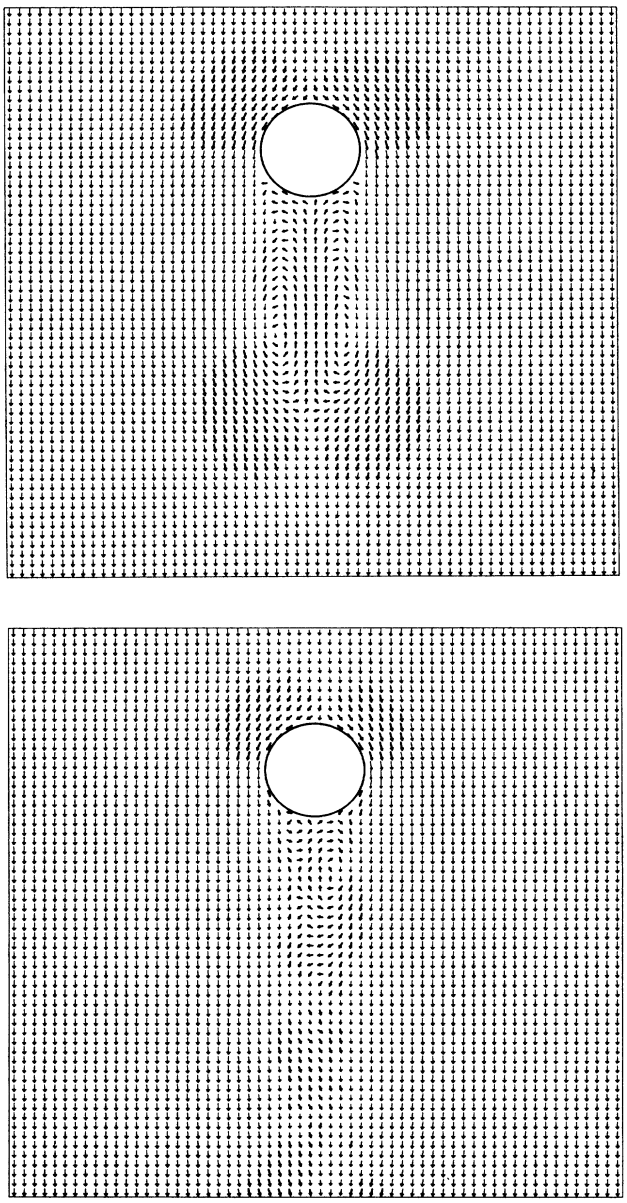

FIG. 5. Arrow plots for run No. 5 showing an elongated eddy pair at time 1050, and at time 3000 an oscillatory wake preceded by an asymmetric pair of eddies.

Taken together, the results suggest that flow instability appears at a nonzero $\mathcal{R}_c$, and that there is a range of $\mathcal{R}$ over which qualitatively similar wake oscillations follow a transient eddy formation stage. Whether the $\mathcal{R}$ range should be further divided because persistent stationary eddies can exist remains a matter for future study. There is also not a great deal of flexibility in the selection of D and U to allow testing the validity of dynamic similarity, namely, that the flow is a function of $\mathcal{R}$ alone. The limitations are that U be sufficiently small to prevent the flow patterns being unduly affected by the boundaries and also that the flow be subsonic, while on the other hand D must also be sufficiently small (relative to the width of the region) to avoid excessive velocities to the sides of the obstacle. Greater flexibility might be provided by obstacles with a lower $\mathcal{R}_c$ than the circle.

Computational power stands in the way of increasing the system size by simply adding particles. A possible alternative route might be to reduce the density. There is obviously a limit to what can be gained by this tactic since the particle mean free path also increases, and one

of the criteria for the applicability of continuum hydrodynamics (also at MD length scales) is that the ratio of mean free path to the characteristic length, in this case D, a quantity known as the Knudsen number ($\mathcal{K}$), be small. At the liquid densities used here $\mathcal{K} \approx 10^{-3}$ because the mean free path is typically 10% of the particle diameter, but an expansion of the system would lead to significantly larger $\mathcal{K}$ where ballistic motion of the particles over distances comparable to D signals the breakdown of hydrodynamics.

One run was carried out with the density reduced by a factor of 16 and the value of D doubled, corresponding to $\mathcal{K} \approx 0.02$. There was no need for as many particles as previously and only 4×10^4 were used, but the region length was still double the previous value. No deviation from Stokes flow was seen. Two explanations of this result, not necessarily mutually exclusive, are that hydrodynamics is no longer relevant and that $\mathcal{R}$ has dropped below $\mathcal{R}_c$ because at low density $\mathcal{R} \propto \mathcal{K}^{-1}$ (an immediate consequence of ν being proportional to the mean free path at low density[12]).

There has been one other MD study with a similar goal.[15] This purportedly dealt with a three-dimensional system, and addressed the flow past a thin inclined plate in the context of rarified-gas dynamics. In actual fact the mean free path was close to the particle size, corresponding to an extremely dense gas, and the extent of the system in one direction was so short as to make the system effectively two dimensional. Vortex shedding was observed in a fluid of 4×10^4 hard spheres (effectively disks). The value of $\mathcal{R}$ was claimed to be 78 but this was based on the dilute-gas assumption; given the actual flow parameters the correct value should be less than that of the present work. Since, as mentioned earlier, $\mathcal{R}_c$ for a sharp-tipped plate is practically zero, the presence of vortices cannot be used to gauge the value of $\mathcal{R}$.

VI. DISCUSSION

The results presented in this paper leave little doubt that physically relevant hydrodynamical instabilities are accessible to simulation at the discrete particle level. With the growing availability of powerful computers it is natural to expect to see more work along the lines of the present study, as well as attempts to explore other fluid flow and stability problems. The issues which need to be addressed, however, including those that might be raised by the devil's advocate, concern the link between microscale phenomena and hydrodynamics at the macroscopic level, possible limitations as to the kinds of effects that are plausible subjects for study, and the prospects of the MD method as a computational tool for fluid engineering.

There is no inherent size limitation built into the foundation of continuum hydrodynamics beyond the requirement that densities and currents exist, and that fluctuations associated with these quantities be negligible. Much of hydrodynamics is devoted to incompressible flow for which the requirement is a low Mach number; transport coefficients (e.g., viscosity) may be regarded as constants or as specified functions of the flow, and temperature gradients excluded. To what extent can features such as these be reproduced in an MD simulation, and how central are they to obtaining meaningful results?

The systems used in the present work are sufficiently large that the coarse-grained flow patterns are both smooth, indicating sufficient particles per cell, and intricate, indicating an adequately small cell. Temperature (or pressure) patterns would be subject to much greater measurement error, especially in regions of complex flow, due to the problem of isolating the thermal component of the particle velocities. In order to achieve a sufficiently high Reynolds number the small size of the obstacle must be compensated by a high flow speed (contrast the slow experimental flows). Despite the flow being subsonic, the compressibility effects are all too apparent, especially the density drop in the shadow of the obstacle. Deviations from incompressible flow would be less significant at reduced flow rates, but this in turn demands a larger obstacle in order to keep $\mathcal{R}$ unchanged, and hence a bigger system. Estimation of $\mathcal{R}$ is problematic given the difficulty in determining the viscosity (see Sec. V); only at sufficiently low flow rates can the viscosity be computed directly from the drag via Stokes' law.

Among better-known flow instabilities that ought to have analogues in two dimensions are the Rayleigh-Bénard and the Kelvin-Helmholtz. As with obstructed flow there are lower limits to system size and observation time inherent in these problems, and these establish the feasibility or otherwise of MD simulation. There are other instabilities that are three dimensional in nature, such as flow past a sphere. To produce a three-dimensional (3D) system whose linear dimension is similar to that of the present 2D system requires at least a hundredfold increase in the number of particles, a problem which may have to await the advances of the next generation of computers. The question of observation time is particularly relevant for the Rayleigh-Bénard instability because, even if the critical Rayleigh number for convection is exceeded, the slow convective velocity associated with the rolls may prove difficult to detect over the thermal background noise, unless the measurement extends over a long period of time.

As to the question whether MD will ever become an engineering tool routinely used in studies of flow problems, it is clearly too early to assess its prospects. It is obvious that the use of MD is totally unnecessary for the majority of flow problems handled perfectly satisfactorily by hydrodynamics. But for flows under extreme conditions, or for a study of what really happens inside a boundary layer, MD may prove to be the only viable approach.

An aspect of the work which has only been briefly alluded to is that of computational efficiency, an important factor in the lengthy calculations necessary for MD simulation. A poorly designed algorithm, inefficient code, or even insufficient awareness of the architectural constraints of the computer can all contribute to the most powerful of computers performing extremely poorly. There is presumably an optimal MD code for each

kind of processor organization, and these may have little in common, but until a machine attuned to the idiosyncracies of the MD method becomes available,[21] a lot can be gained by adjusting the computational approach to the hardware used.

VII. FUTURE PROSPECTS

Scattered throughout the paper, both implicitly and explicitly, are suggestions as to how the present exploratory work might be extended given an adequate allocation of time on a "supercomputer" (defined, incidentally, to be a system that falls only one generation behind the computing requirements of leading-edge efforts in science and engineering[35]). A number of these are summarized here.

A great deal can be done without going to even larger systems. Perhaps the two most interesting questions are (i) whether it is feasible to produce a stationary pair of eddies and (ii) the effect of altering the boundary conditions from slip to nonslip (a cursory examination suggests that the eddies are enhanced—as might be expected). In order to systematically study different flow rates it is desirable to incorporate a feedback mechanism that adjusts the field strength in response to flow velocity changes; such a "velostat" is complicated by the relatively long reaction time to a variation in field strength (typically the traversal time of the system) and the fact that the drag exerted by the obstacle is both a function of $\mathcal{R}$ and subject to fluctuation. Other obstacle shapes can be explored, in some cases even using smaller systems, as can the effect of density changes. Where feasible, the macroscopic observation that flow depends on $\mathcal{R}$ alone should be tested at the MD level. Achieving significantly higher Reynolds numbers, or simulating systems of corresponding size in three dimensions, may have to wait for a later generation of computers.

Only a single form of potential function was used throughout the simulations. An attractive tail could be added to the interaction, thereby recovering the original Lennard-Jones potential, in order to see whether this reduces the density change in the lee of the obstacle by making the fluid more resistant to cavitation—the repulsive potential has no such resistance. Then, of course, there are a host of other fluids—non-Newtonian, reacting mixtures, polymers and suspensions. For many problems MD simulation may prove to be the only viable approach in the long term.

ACKNOWLEDGMENTS

The computations described here were carried out while visiting IBM in Kingston, N.Y. The author wishes to thank Dr. E. Clementi for his hospitality.

[1]G. K. Batchelor, *An Introduction to Fluid Dynamics* (Cambridge University, Cambridge, 1967).
[2]D. C. Rapaport and E. Clementi, Phys. Rev. Lett. **57**, 695 (1986).
[3]P. W. Bearman and J. M. R. Graham, J. Fluid. Mech. **99**, 225 (1980).
[4]J. H. Gerrard, Philos. Trans. R. Soc. London Ser. A **289**, 351 (1978).
[5]M. van Dyke, *An Album of Fluid Motion* (Parabolic, Stanford, CA, 1982).
[6]A. E. Perry, M. S. Chong, and T. T. Lim, J. Fluid Mech. **116**, 77 (1982).
[7]K. Clark and C. Pedretti, *The Drawings of Leonardo da Vinci in the Collection of Her Majesty the Queen at Windsor Castle,* 2nd ed. (Phaidon, London, 1968).
[8]J. W. S. Rayleigh, *The Theory of Sound* (Dover, New York, 1945).
[9]D. C. Thoman and A. A. Szewczyk, Phys. Fluids Suppl. **II**, 76 (1969).
[10]S. K. Jordan and J. E. Fromm, Phys. Fluids **15**, 371 (1972).
[11]N. Takemitsu, J. Phys. Soc. Jpn. **48**, 306 (1980).
[12]L. D. Landau and E. M. Lifshitz, *Fluid Mechanics* (Pergamon, Oxford, 1959).
[13]D. S. Dandy and L. G. Leal, Phys. Fluids **29**, 1360 (1986).
[14]G. A. Bird, *Molecular Gas Dynamics* (Clarendon, Oxford, 1976).
[15]E. Meiburg, Phys. Fluids **29**, 3107 (1986).
[16]U. Frisch, B. Hasslacher, and Y. Pomeau, Phys. Rev. Lett. **56**, 1505 (1986)
[17]N. Margolis, T. Toffoli, and G. Vishniac, Phys. Rev. Lett. **56**, 1694 (1986).
[18]J. B. Salem and S. Wolfram, Thinking Machines Corp. (Cambridge, MA), Technical Report No. 86.14, 1986 (unpublished).
[19]D. d'Humieres, Y. Pomeau, and P. Lallemand, C. R. Acad. Sci., Ser. B **301**, 1391 (1985).
[20]*Molecular Dynamics Simulation of Statistical Mechanical Systems,* Proceedings of the International School of Physics "Enrico Fermi," Course XCVII, Varenna, 1985, edited by G. Ciccotti and W. G. Hoover, (North-Holland, Amsterdam, 1986).
[21]F. F. Abraham, Adv. Phys. **35**, 1 (1986).
[22]D. J. Evans and W. G. Hoover, Annu. Rev. Fluid Mech. **18**, 243 (1986).
[23]J. J. Erpenbeck and W. W. Wood, in *Modern Theoretical Chemistry,* edited by B. J. Berne (Plenum, New York, 1977), Vol. 6B, p. 1.
[24]D. C. Rapaport, J. Comput. Phys. **34**, 184 (1980).
[25]C. W. Gear, *Numerical Initial Value Problems in Ordinary Differential Equations* (Prentice-Hall, Englewood Cliffs, NJ, 1971).
[26]B. B. Mandelbrot, *The Fractal Geometry of Nature* (Freeman, San Fransisco, 1982).
[27]D. J. Evans and G. P. Morriss, Comput. Phys. Rep. **1**, 297 (1984).
[28]D. J. Evans and G. P. Morriss, Phys. Rev. Lett. **56**, 2172 (1986).
[29]R. W. Hockney and C. R. Jesshope, *Parallel Computers* (Hilger, Bristol, 1981).
[30]G. C. Fox and S. W. Otto, Phys. Today **39** (5), 50 (1984).
[31]E. Clementi and D. Logan, IBM Kingston Report No. KGN-43, 1985 (unpublished).
[32]D. Levesque, L. Verlet, and J. Kurkijarvi, Phys. Rev. A **7**, 1690 (1973).
[33]H. Schlichting, *Boundary-Layer Theory,* 6th ed. (McGraw-Hill, New York, 1968).
[34]D. M. Heyes, G. P. Morriss, and D. J. Evans, J. Chem. Phys. **83**, 4760 (1985).
[35]N. R. Lincoln, Computer **16** (5), 38 (1983).