Molecular-dynamics simulation of compressible fluid flow in two-dimensional channels

M. Sun and C. Ebner
Department of Physics, The Ohio State University, Columbus, Ohio 43210
(Received 24 April 1992)

We study compressible fluid flow in narrow two-dimensional channels using a molecular-dynamics simulation method. In the simulation area, an upstream source is maintained at constant density and temperature while a downstream reservoir is kept at vacuum. The channel is sufficiently long in the direction of the flow that the finite length has little effect on the properties of the fluid in the central region. The simulated system is represented by an efficient data structure, whose internal elements are created and manipulated dynamically in a layered fashion. Consequently the computer code is highly efficient and manifests completely linear performance in simulations of large systems. We obtain the steady-state velocity, temperature, and density distributions in the system. The velocity distribution across the channel is very nearly a quadratic function of the distance from the center of the channel and reveals velocity slip at the boundaries; the temperature distribution is only approximately a quartic function of this distance from the center to the channel. The density distribution across the channel is nonuniform. We attribute this nonuniformity to the relatively high Mach number, approximately 0.5, in the fluid flow. An equation for the density distribution based on simple compressibility arguments is proposed; its predictions agree well with the simulation results. The validity of the concept of local dynamic temperature and the variation of the temperature along the channel are discussed.
PACS number(s): 47.40.Dc, 47.60.+i

I. INTRODUCTION

The technique of molecular-dynamics (MD) simulation has been widely used to study nonequilibrium fluids. Because of limitations imposed by finite computational capacity with regard to both memory and speed, this method has been used principally to determine the behaviors of fluid systems on time and distance scales within a few orders of magnitude of τ_c and r_c, respectively. Here, $\tau_c \approx 10^{-13}$ sec is the collision duration and $r_c \approx 10^{-8}$ cm is a molecular size [1]. Only molecular properties of the fluid can be obtained in this range of time and distance, and that is also the domain of experimental neutron-scattering measurements. Realistic examination of the hydrodynamic properties of flow is still beyond the reach of most molecular-dynamics simulations, although there are many attempts to simulate larger systems on longer time scales. Among these are the work of Koplik, BanxDavar, and Willemsen [2,3], Hannon, Lie, and Clementf [4], and Bhattacharya and Lie [5, 6]. In particular, Hannon, Lie, and Clementi have obtained velocity and temperature distributions across channels in which flow is occurring. Their results agree well with simple hydrodynamic predictions for incompressible fluid flow. Also, Bhattacharya and Lie have obtained similar velocity profiles and have computed boundary-slip coefficients.

These simulations of fluid flow have a common property. Periodic boundary conditions are introduced along the flow direction in the interest of reducing the amount of computation needed to obtain useful results. The imposition of translational invariance along the flow direction makes it necessary to introduce a "gravitational" field to induce flow. In order to induce appreciable flow, this field must be given a strength much larger than the earth's field g, for example, as large as [2] $10^{12}g$. As a consequence of the gravitational field, regular rescaling of the particles' kinetic energies is required; hence one cannot reliably study properties of the fluid having to do with energy or heat flow. Also, it is not possible to study variations of the flow properties along the flow direction.

In this paper, we present a method of simulation in which the channel is of finite length in the flow direction without periodic boundary conditions. At one end of the channel is a source region which is maintained at constant density and temperature by introducing particles as needed. At the other end is a sink region which is maintained at vacuum; that is, any particle which moves into this region is removed from the system. Hence the pressure or density gradient along the channel is primarily responsible for instigation of the flow. We also use a layered data structure to improve the efficiency of the computer code and the utilization of storage. Our methods make it possible to simulate systems containing 20 000 or more particles for more than 10^6 time steps on a DECstation 3100 computer.

Section II contains a description of the system simulated and of our numerical techniques. The results are in Sec. III, and Sec. IV contains a discussion and conclusions.

II. MODEL AND NUMERICAL METHODS

The geometry of the system is indicated in Fig. 1. The channel has a length L in the x direction and width w

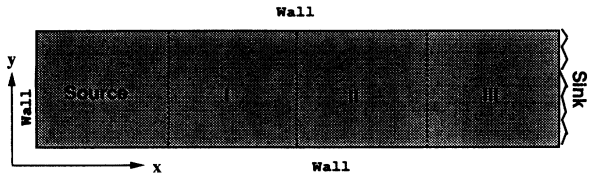

FIG. 1. Diagram of the two-dimensional channel showing the source and sink regions. Walls are indicated by solid lines. Average cross-channel distributions presented in some of the following figures are obtained in region II which is equidistant from the two ends of the simulation area, excluding the source region. The lengths of the channel in the x and y directions are denoted in the text by L and w, respectively; the length of the source in the x direction is called L_1.

in the y direction. Flow is along the x direction with the source region at $0 < x < L_1$ and the vacuum or sink located at $x > L$. The channel is closed at $x = 0$ so that particles cannot escape into the region $x < 0$. Typical channel sizes that we have investigated are $L = 400\sigma$, $w = 100\sigma$, and $L_1 = 100\sigma$, where σ is the particle-size parameter (diameter) in the Lennard-Jones potential

$$V(r) = 4\varepsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right], \tag{1}$$

which we have employed for the interparticle interactions. The interaction is truncated at $r \geq 2\sigma$ [7]. In addition to ε and σ, the only other parameter describing the properties of the molecules is the mass m. We have used values appropriate for argon, i.e., $\sigma = 3.4\,\text{Å}$, $\varepsilon/k = 119.76\,\text{K}$, and $m = 6.67 \times 10^{-23}\,g$; k is the Boltzmann constant. We also introduce a basic time constant $\tau = \sqrt{m\sigma^2/48\varepsilon} \approx 3 \times 10^{-13}\,\text{sec}$.

Our procedure for handling collisions of the particles with the walls, which are the surfaces $x = 0$, $y = 0$, and $y = w$, is to give the recoiling particles a Maxwell-Boltzmann velocity distribution on a half-space. For example, after colliding with the wall at $x = 0$, a particle is given a velocity $\mathbf{v} = v_x\hat{\mathbf{x}} + v_y\hat{\mathbf{y}}$ (where $v_x > 0$) with a probability proportional to $\exp[-m(v_x^2+v_y^2)/2kT]$ where T is the specified wall temperature. Certainly, this model for the particle-wall collisions is only a crude approximation to what must happen at a real wall. It does allow for energy transfer to and from the fluid and will maintain a given temperature or mean kinetic energy in the system. Undoubtedly, the microscopic properties of the fluid, such as the velocity distribution of the particles, close to the wall are strongly affected by the wall-particle collisions. In particular, the behavior of the fluid temperature and drift velocity in this region and the energy transfer between fluid and wall are suspect. This point seems to us sufficiently important that we have embarked on further simulations incorporating a different, presumably more realistic, treatment of the wall-fluid interface.

The molecular-dynamics simulation proceeds in the conventional manner except that when the particle density in the source region drops by a small amount, typically after several hundred time steps, we inject particles into this region in order to bring the density here back up to some prescribed value. That is done by, first, letting the end wall ($x = 0$) act as a piston and uniformly compress the particles in the source region (but not the remainder of the channel) by a sufficient amount to bring the density in this compressed volume up to the desired level, and, second, restoring the end wall to its original position and injecting an appropriate number of particles into the empty space next to the wall. These added particles are given a Maxwell velocity distribution at the same temperature T as that of the walls. By making this adjustment to the number of particles sufficiently frequently, we only have to inject a few particles each time and the relative amount by which the source region is compressed is very small. By comparison, the entire system typically contains 10 000–20 000 particles, and the source region is one-fourth of the channel.

The initial state of the system is always taken to have some preset density of particles with a Maxwell-Boltzmann velocity distribution in the source region and no additional particles elsewhere. Thus, as time proceeds, particles make their way down the channel, driven by the pressure or density gradient, and are removed at the far end. After a sufficiently long time a steady state is achieved and various quantities such as the distributions of flow velocity, particle density, particle-current density, and temperature, may be computed. We typically run the simulation in the steady state for some 10^6 time steps with one time step being $\delta t = 3 \times 10^{-2}\tau \approx 9 \times 10^{-15}$ sec. In addition, some 10^6 time steps are needed for the system to reach the steady state.

The most critical part of a molecular-dynamics simulation is the force calculation [8] which would require $\sim N^2$ steps in a brute force calculation if N is the number of particles. Given short-ranged interactions, such as the Lennard-Jones potential, one typically uses Verlet's neighbor list[8] to lower the number of steps to $\sim NN_c$, where N_c is the maximum number of particles that can fit within the range of a particle's potential. This method has high memory requirements and also requires a number of steps $\sim N^2$ to update the neighbor list. A superior method [8] in large systems is to divide the simulation area into cells and use two arrays HEAD and LIST to store linked lists. This method uses no extra memory and has linear performance, i.e., $O(N)$ steps. However, the reference locality is broken when molecules in a given cell are stored in a long array LIST for a large system. This can be detrimental to the performance of the computer code.

In our implementation of the MD simulation we build a specially tailored data structure to provide a better mapping between the physical problem and the memory partition of modern computers. It is shown schematically in Fig. 2. A node, which is roughly speaking a piece of computer memory, is created for every particle introduced into the system. All the information for that particle is stored in the node. The memory of the node is deallocated when the particle moves out of the system. The system is divided into cells to reflect the short-ranged nature of the interaction. The neighboring environment of every molecule is implemented by a dynamic one-way link list in each cell. This one-to-one

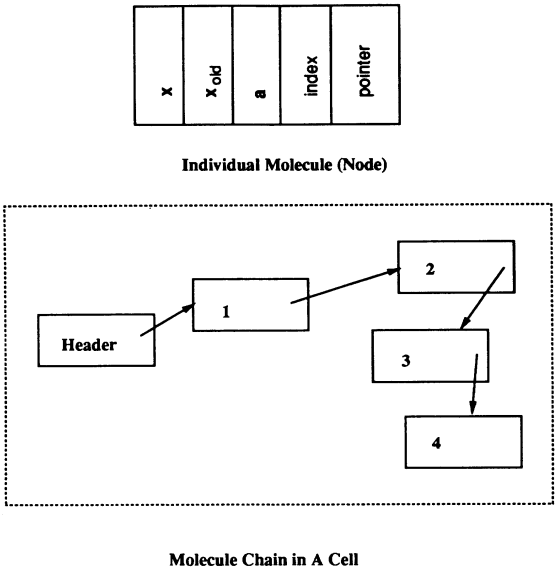

FIG. 2. Schematic representation of the data structures employed in the molecular-dynamics simulations. Each node stores all information such as position, acceleration, etc. of a particle. The simulation area is divided into rectangular cells. Particles (nodes) in each cell are linked by pointers to form a dynamic link list.

dynamic management of the simulated systems enables a high degree of reference locality and efficient usage of computer resources.

This data representation also enables layered structures. We already see the atom-node layer and the cell layer. Further layers can be built on top of these to take into account the structures of, e.g., polyatomic molecules or other complex objects. Hence this method of constructing the data improves the code reusability and makes simulation of complex systems relatively easy. We have, for example, extended our code to simulate the flow of a fluid composed of propane (C_3H_8) molecules. With this data structure, we are able to achieve on a single DECstation 3100 computer a speed of 3×10^{-5} sec/particle per step, comparable or superior to that obtained in some [3, 9] but not all [10] simulations done on supercomputers for comparable systems.

III. RESULTS

All of the figures in this section are for steady-state flow with $L = 400\sigma$, and $w = 100\sigma = L_1$; in the source region, $T = \varepsilon/k$ and $n = 0.25/\sigma^2$. The lone exception is the inset in Fig. 5 which is for a narrow channel with $w = 40\sigma$. Also, the density at small x is larger than 0.25 in this narrow channel. The wall and source region temperature is well above the critical temperature $T_c \approx 0.55\varepsilon/k$ of the two-dimensional Lennard-Jones fluid [11] and the density is well below solid densities $n_s \geq 0.8/\sigma^2$ so that we will have a one-phase system unless sufficient cooling takes place at some point along the channel to produce a much lower temperature. Some cooling is expected as the fluid expands and moves down the channel; however, collisions with the walls, which are also at $T = \varepsilon/k$, will compensate somewhat for this effect.

Figure 3 shows the particle-current density along the channel, J_x, in units of $1/(\sigma\tau)$, as a function of x/σ. Different sets of points correspond to different intervals of y. For example, the points labeled "0–10" represent an average of the current density for $0 < y < 10\sigma$, and of the current density for $w - 10\sigma < y < w$. Because the distribution is symmetric around the line $y = w/2$, we can average over the data from strips symmetrically placed on the two sides of this line. We have done so in order to obtain smoother results. The current density averaged across the entire channel is also shown. The constancy of this density outside of the source region, as a function of x, is an indication that the system is indeed in the steady state. The data in this figure, as well as in all of the following ones, are obtained by averaging over a million time steps after discarding results from an initial one million steps during which the steady state is achieved.

Two other features of J_x are worthy of comment. First, the current density increases approximately linearly with x in the source region reflecting the fact that there is a hard wall at $x = 0$ where J_x must vanish. Second, the current densities in different intervals in the y direction are not all independent of x in the middle ($100\sigma < x < 300\sigma$) of the channel but show roughly linear variation. Moreover, this variation is such that J_x increases (decreases) along the pipe far from (close to) the walls. Correspondingly, we find that at the smaller values of x, there is a flow of particles, described by J_y, toward the center of the channel while at the other end there is a flow toward the sides. This behavior is present

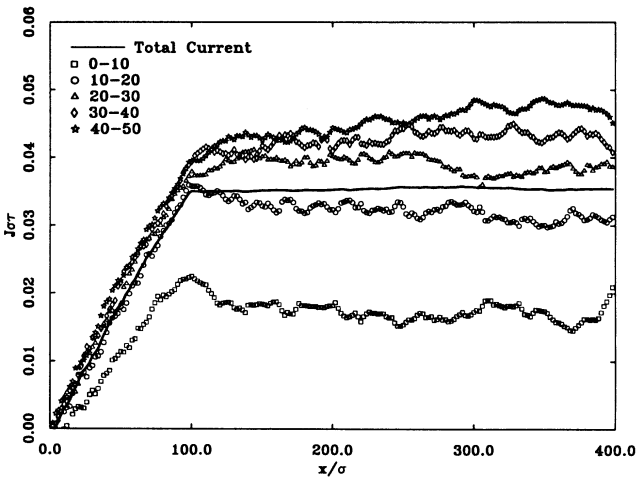

FIG. 3. The particle-current density J_x, in units of $1/\sigma\tau$ and averaged within various intervals $\Delta y/w$ across the channel as indicated by the legends, is shown as a function of x/σ. Also shown as a solid line is J_x averaged across the entire channel. These results are for a channel with $w = 100\sigma$ and $L = 400\sigma$; the source and wall temperature is $T = \varepsilon/k$; and the averages are taken over 10^6 time steps after an initial 10^6 steps during which the system reaches the steady state.

because the system is not a great deal longer than wide. In a channel with the same length but a width of only 40σ, the variation of J_x with x is relatively much less as is the magnitude of the transverse current density J_y.

In Fig. 4 we show the particles' mean velocity in the x direction, u, in units of σ/τ, as a function of the position across the channel y/w. The velocity is evaluated in region II of the channel which is the middle third of the flow area excluding the source region. In this part of the channel the end effects are minimized. The velocity u is the macroscopic fluid velocity. One may see clearly the velocity slip at the walls. Also given in the figure as a solid line is a parabola which has been fit to the data. For an incompressible fluid and laminar flow, the velocity profile is expected to be parabolic. From the figure one can see that, even though we have a highly compressible fluid, cf. Figs. 5 and 6, the distribution still fits the parabola quite well.

The cross-channel density distribution in region II is shown in Fig. 5, and the density along the channel, averaged across the channel, appears in Fig. 6. One can see that n decreases roughly, but definitely not precisely, linearly as x increases. Also, it is significantly lower in the middle of the channel than toward the sides. The cross-channel variation is expected for compressible fluid flow when the velocity is large enough [12]. For Mach number less than about 0.3, the density variation is usually not great but when it is around 0.5 or larger, which is the case in our system, the finite compressibility of the fluid is highly evident. Exact solutions of the hydrodynamic equations for the flow of a compressible fluid are known only for a few special cases. We propose the following simple arguments to explain qualitatively the density distribution observed in our simulations. First, since the local temperature variation across the channel is quite small in comparison with the total kinetic-energy variation, we shall neglect the former. Then to lowest order in the Mach number M we have the following generalized Bernoulli equation for isentropic flow:

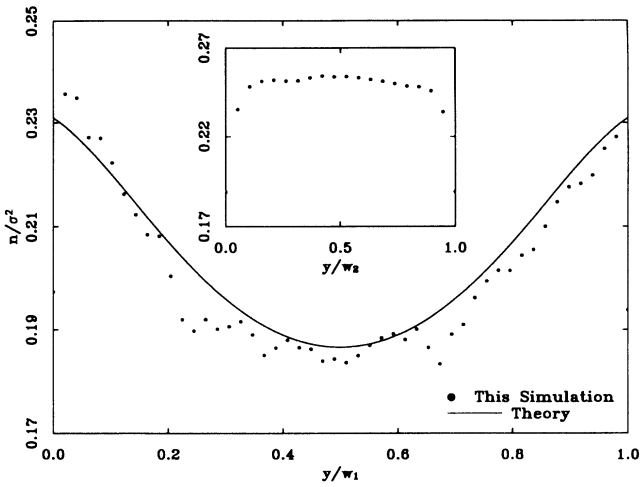

FIG. 5. The steady-state cross-channel particle density n in region II in units of $1/\sigma^2$ is shown as a function of x/σ for a wide channel ($w = 100\sigma$) as a function of y/w. The solid line is the theory described in the text. The inset displays the particle density for a narrow channel ($w = 40\sigma$). Other parameters are the same as for Fig. 3 except that the maximum density in the narrow channel is maintained at about $0.28/\sigma^2$.

$$P_0 - P = \tfrac{1}{2}\rho c^2 M^2, \tag{2}$$

where c is the speed of sound; P is the pressure at a point where the flow velocity is u; and P_0 is the pressure where $u = 0$, or stagnation pressure. The speed of sound can be approximated as follows:

$$c^2 = \left(\frac{\partial P}{\partial \rho}\right)_S \approx \left(\frac{\Delta P}{\Delta \rho}\right)_S, \tag{3}$$

and we can identify $\Delta P = P - P_0$ and $\Delta\rho = \rho - \rho_0$

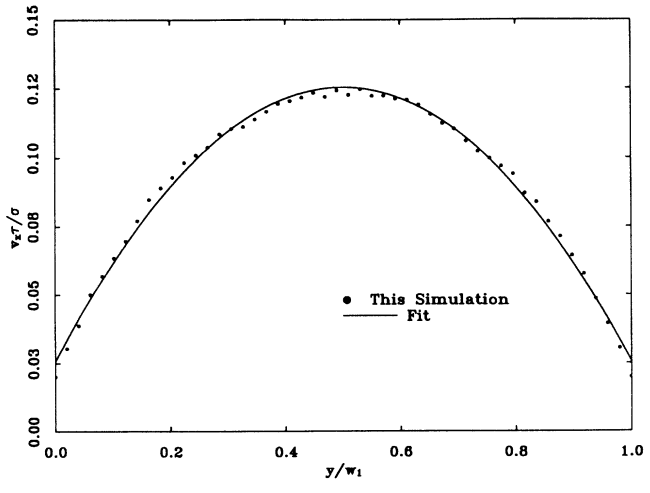

FIG. 4. The steady-state average fluid velocity u in region II is plotted against y/w; u is given in units of σ/τ. The parameters are the same as for Fig. 3. The solid line represents the best fit of a parabola to the simulation results.

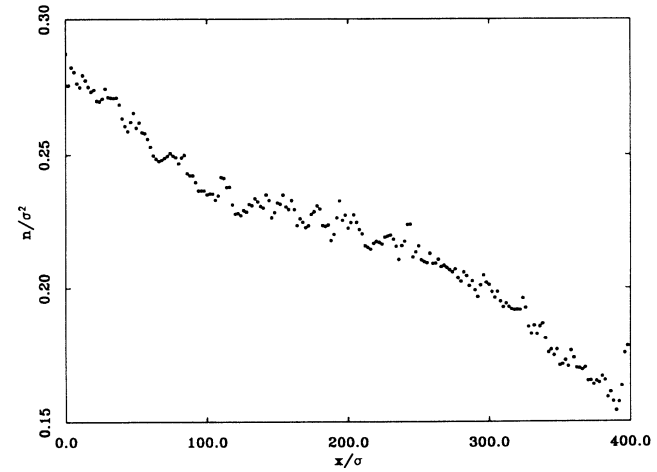

FIG. 6. The steady-state average density along the channel in units of $1/\sigma^2$ is shown as a function of x/σ. Parameters are the same as in Fig. 3.

where ρ_0 is the stagnation density. Then we have, still to second order in M,

$$\rho = \frac{\rho_0}{1+\frac{1}{2}M^2} = \frac{\rho_0}{1+\frac{1}{2}\frac{u^2}{c^2}}. \quad (4)$$

This function is plotted as a solid line in Fig. 5 with u computed from the simulation. For ρ_0 we use the density at the downstream end of the source region, $\rho_0 = 0.23/\sigma^2$. We did not use the mean density in this region because of the unphysical rescaling we use here in the simulations to balance the extra heat generated by bringing in new particles. In addition, the speed of sound is temperature dependent and is not a constant throughout the channel leading to a further ambiguity in connection with the appropriate choice of c. If we simply choose c to be the ideal-gas value $c = \sqrt{2kT/m}$ with $T = 0.75/k$, which is the average temperature in region II of the channel, cf. Fig. 8, that gives the result shown in Fig. 5. Qualitatively, the fit is quite reasonable, perhaps better than one might have expected.

Also shown as an inset in Fig. 5 is the cross-channel density distribution in region II for a considerably narrower channel with $w = 40\sigma$. In this case the density is quite constant across the channel except very close to the edges where it decreases somewhat. By contrast, the cross-channel velocity distribution for the narrow channel shows behavior very similar to that for the wide channel except that the peak velocity in the center is considerably smaller in the former case. The flow is sufficiently slow (the Mach number is around 0.25 midway along the channel) that at a given x no density decrease is apparent in the middle of the channel as compared with the density closer to the sides.

For both the narrow channel and the wide one, there is a significant decrease in the particle density very close to the walls, i.e., within about 2σ of the walls. We believe this decrease to depend on, and to be a consequence of, the manner in which we reflect particles from a wall.

The local temperature distribution in the channel has been obtained from study of the velocity distributions of the particles within each cell. Of course, one may ask whether there is a well-defined temperature in an intrinsically nonequilibrium system. We obtain the temperature from the equipartition theorem according to which, in the absence of any flow, $kT = m\overline{v_x^2}$ or $kT = m\overline{v_y^2}$ where the overlines indicated mean values taken over the distribution of particle velocities; both temperatures should be the same. However, given that there is some flow velocity u in the x direction as in our channel, then if there is a Maxwell velocity distribution centered at this velocity, we may introduce temperatures

$$kT_x = m\overline{(v_x-u)^2}, \quad kT_y = m\overline{v_y^2}. \quad (5)$$

Using these formulas we have computed T_x and T_y. The cross-channel temperature distributions in region II, in units of ε/k, are shown in Fig. 7. The two generally agree quite well except very close to the edges of the channel where T_x is considerably larger than T_y; specifically, $T_x \approx 0.76$ and $T_y \approx 0.67$ right at the wall. One should expect that the two would not agree here because the velocity distribution is certainly not a moving Maxwell distribution as a consequence of the particular manner in which we handle reflections of the particles from the walls. Rather, the distribution close to a wall is a superposition of two distributions; the reflected particles have a Maxwell distribution which is not moving while the incident particles have a net velocity along the x direction. The solid curve in Fig. 7 is the best fit of a quartic function to the data. A quartic has been used because that is the predicted temperature distribution for an incompressible fluid. One can see that the fit is not very good, indicating that the prediction of incompressible fluid theory is really not appropriate in this case. In addition to the fact that our fluid is manifestly compressible, the nature of the wall-fluid interaction may be a contributing factor to deviations of the temperature profile from quartic behavior close to the walls.

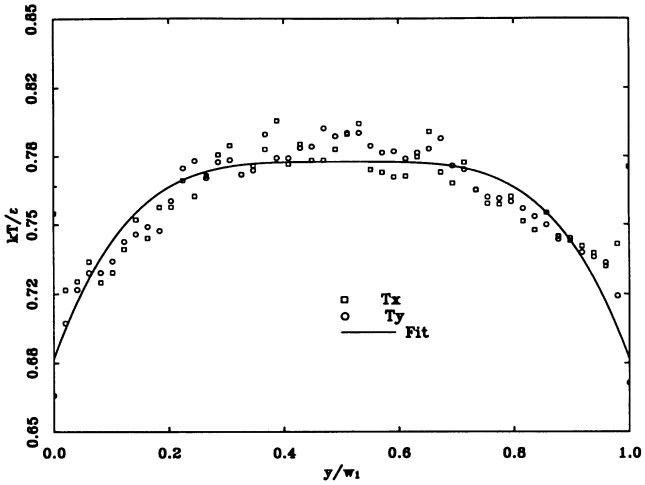

FIG. 7. The steady-state cross-channel temperature distributions T_x and T_y in region II, extracted from the equipartition theorem by computing the particles' kinetic energies, are given in units of ε/k as functions of y/w. Parameters are the same as for Fig. 3. The solid curve is the result of a least-squares fit of a quartic function to the simulation results.

Figure 8 shows the average temperature distribution along the channel. One can see that T_x and T_y agree quite well and that the temperature drop, once the particles have left the source region, is roughly linear. Several factors contribute to the temperature drop. First, the fluid is expanding, which will increase its potential energy, and so to the extent that energy is conserved the mean kinetic energy per particle must decrease. Second, collisions with the walls offset this effect because particles reflected from the walls have on the average a kinetic energy of kT where the wall temperature T is the same as the initial temperature of particles injected into the source region. Finally, there is also the effect of the increasing flow velocity along the channel which has the consequence that, as x increases, more of the kinetic energy is going into the flow and correspondingly less energy is available for fluctuations of the particles' velocities relative to the flow velocity, thereby producing

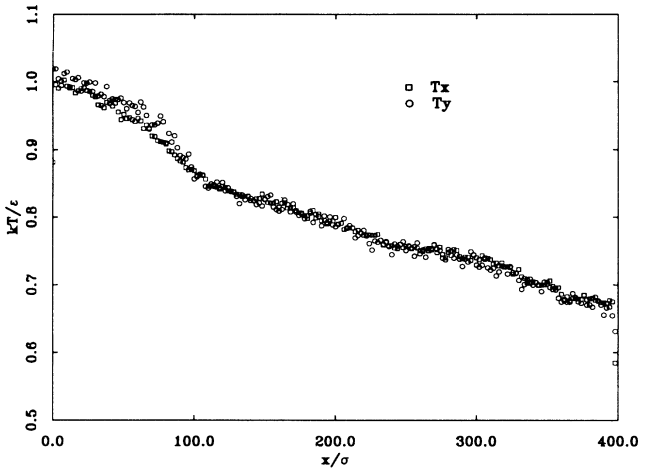

FIG. 8. The steady-state temperatures T_x and T_y in units of ε/k, averaged across the channel, are shown as functions of x/σ. Parameters are the same as for Fig. 3.

a drop in the temperature, as we have defined it, with increasing x.

IV. CONCLUSIONS

We find that, in two dimensions at least, it is practical to simulate a realistic flow system with a source and a drain, and at the same time keep a long enough channel that finite-size effects associated with the length are not important in regions farthest from the source or drain. Sufficiently high Mach numbers to allow for the investigation of compressible fluid flow can be achieved as is evidenced by the considerable cross-channel density variations observed. The nonuniformity of this density distribution can be roughly explained by an equation based on the Bernoulli equation and the relation between system compressibility and the sound velocity. This equation also predicts, and our results verify, that the density distribution will be more nearly uniform (aside from the region within a few mean free paths of the walls) for flow with small Mach numbers.

We also find that there is good agreement between our results for the flow velocity and predictions based on the hydrodynamics of incompressible fluids. Overall, the results suggest that for fairly high ($M \sim \frac{1}{2}$), but subsonic, Mach number flow, a fairly good picture of the system's behavior can be obtained from solutions for incompressible fluid flow with some corrections of order M^2 to account for the consequences of compressibility.

ACKNOWLEDGMENTS

We wish to thank Gary Grest for a useful discussion concerning molecular-dynamics simulations of fluids and for bringing Ref. 10 to our attention. This work was supported by National Science Foundation Grant No. DMR-9014679.

[1] J. P. Boon and S. Yip, *Molecular Hydrodynamics* (McGraw-Hill, New York, 1980).
[2] J. Koplik, J. R. BanxDavar, and J. F. Willemsen, Phys. Fluids A **1**, 781 (1989).
[3] J. Koplik, J. R. BanxDavar, and J. F. Willemsen, Phys. Rev. Lett. **60**, 1282 (1988).
[4] L. Hannon, G. C. Lie, and E. Clementi, Phys. Lett. A **119**, 174 (1986).
[5] D. K. Bhattacharya and G. C. Lie, Phys. Rev. Lett. **62**, 897 (1989).
[6] D. K. Bhattacharya and G. C. Lie, Phys. Rev. A **43**, 761 (1991).
[7] For every particle in a cell, we only calculate its interaction with particles in the same cell and neighboring ones. Since each cell has a length of 2σ, the cutoff distance of our force calculation is at least 2σ in worst cases.
[8] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Clarendon, Oxford, 1987).
[9] P. Borgelt, C. Hoheisel, and G. Stell, Phys. Rev. A **42**, 789 (1990).
[10] G. S. Grest, B. Dünweg, and K. Kremer, Comput. Phys. Commun. **55**, 269 (1989).
[11] F. F. Abraham, Phys. Rep. **80**, 340 (1981).
[12] P. A. Thompson, *Compressible-Fluid Dynamics* (McGraw-Hill, New York, 1972).

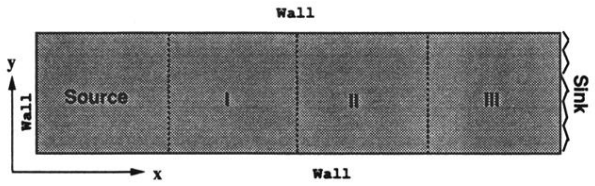

FIG. 1. Diagram of the two-dimensional channel showing the source and sink regions. Walls are indicated by solid lines. Average cross-channel distributions presented in some of the following figures are obtained in region II which is equidistant from the two ends of the simulation area, excluding the source region. The lengths of the channel in the x and y directions are denoted in the text by L and w, respectively; the length of the source in the x direction is called L_1.