MOLECULAR SIMULATION OF RHEOLOGICAL PROPERTIES USING MASSIVELY PARALLEL SUPERCOMPUTERS

B. K. Bhupathiraju, S. T. Cui, S. A. Gupta
Department of Chemical Engineering
University of Tennessee
Knoxville, TN 37996-2200

H. D. Cochran
Chemical Technology Division
Oak Ridge National Laboratory*
Oak Ridge, TN 37831-6224

P. T. Cummings
Department of Chemical Engineering
University of Tennessee
Knoxville, TN 37996-2200

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

Molecular Simulation of Rheological Properties using Massively Parallel Supercomputers

R. K. Bhupathiraju
University of Tennessee
Department of Chemical Engineering
Knoxville, TN 37996-2200
ravib@verdi.engr.utk.edu
http://flory.engr.utk.edu/~ravi

S. T. Cui
University of Tennessee
Department of Chemical Engineering
Knoxville, TN 37996-2200
rcui@utkux1.utk.edu
http://flory.engr.utk.edu/~cui

S. A. Gupta
University of Tennessee
Department of Chemical Engineering
Knoxville, TN 37996-2200
sunil@verdi.engr.utk.edu
http://flory.engr.utk.edu/~sunil

H. D. Cochran
Oak Ridge National Laboratory
Chemical Technology Division
Oak Ridge, TN 37831-6224
hdc@ornl.gov
http://www.ornl.gov/divisions/ctd/EnergyResearch/SolutionThermo/Cochran.html

P. T. Cummings
University of Tennessee
Department of Chemical Engineering
Knoxville, TN 37996-2200
ptc@utk.edu
http://flory.engr.utk.edu/ptc/bio.html

Abstract:

Advances in parallel supercomputing now make possible molecular-based engineering and science calculations that will soon revolutionize many technologies, such as those involving polymers and those involving aqueous electrolytes. We have developed a suite of message-passing codes for classical molecular simulation of such complex fluids and amorphous materials and have completed a number of demonstration calculations of problems of scientific and technological importance with each (described at the World Wide Web site http://flory.engr.utk.edu/ldrd). In this paper, we will focus on the molecular simulation of rheological properties, particularly viscosity, of simple and complex fluids using parallel implementations of non-equilibrium molecular dynamics. Such calculations represent significant challenges computationally because, in order to reduce the thermal noise in the calculated properties within acceptable limits, large systems and/or

long simulated times are required.

1. Introduction

Molecular simulation can be classified into two broad categories: Molecular dynamics simulation and Monte Carlo simulation. In molecular dynamics, Newton's equations of motion for the positions and velocities of the molecular units (which may be atoms, molecules or functional groups within molecules) are solved numerically to determine the properties of interest (such as energy, pressure, stress tensor and measures of structure) over a period of time (typically of the order of $10^5 - 10^8$ time steps) so that the properties can be averaged over time to yield predictions for the macroscopically observable quantities. Since the time step is determined by the need to represent molecular collisions accurately, it is typically of the order of $10^{-15} - 10^{-14}$ seconds, or approximately 1/100th of the mean time between collisions in a liquid. The presence of forces within molecules (intramolecular forces) which are very strong in character - such as the carbon-carbon bond interaction in a hyrdocarbon chain - can in principle require a much smaller time step due to the characteristic relaxation time for such motions. In Monte Carlo simulation, configurations of the system (positions only) are generated randomly by a Markov chain process in order to generate configurations distributed according to the relevant probability distribution (e.g., Boltzmann distribution). Properties of interest are then averaged over the configurations generated. Monte Carlo simulation has the advantage that cleverly constructed Markov chains can explore much wider regions of configurational space than molecular dynamics for an equivalent investment of computing resources. However, unlike molecular dynamics, Monte Carlo cannot be used to study many dynamical phenomena, such as transport properties, because of its inherent lack of dynamical information.

For several decades, molecular simulation has played a major role in physical properties research in providing essentially exact quasi-experimental data for comparison to the results of statistical mechnical theories [1]. The advent of massively parallel supercomputers now makes it possible for molecular simulation methods to become the direct route to predicting the properties of real systems of industrial interest. In particular, the prediction of the properties of bulk fluid-phase and amorphous systems is now within reach, whereas previously molecular simulation was largely restricted to small fluid systems (such as rational drug design, in which 1-2 large biomolecules are suspended in a fluid solvent) or crystalline systems (where the regularity of the structure over very small distance scales renders the problem comparatively straightforward computationally). Using molecular simulation methods implemented efficiently on massively parallel supercomputers, we believe that we are on the threshold of a new era of molecular-based engineering and science, in which the properties of complex liquids and amorphous materials can be confidently predicted using molecular simulation methods, and thus designed and optimized for specific applications prior to synthesis. Just as such approaches have made a significant impact on drug design, one can expect that these methods will have a major impact on the design of, for example, high-performance polymers, synthetic lubricants, and environmentally-benign solvents for chemical reactions and separations.

In this paper, we will focus on the calculation of viscosity using the technique of non-equilibrium molecular dynamics. Both experimentally and computationally, viscosity is typically measured by subjecting a fluid to planar Couette flow. Fluid is placed between two parallel planar walls, one of which moves at a velocity V with respect to the other/ This sets up at steady state a velocity profile across the

wall, since fluid in contact with the moving wall has net velocity V, while fluid in contact with the stationary wall has no net velocity. The force exerted in the x-direction on the upper plate in Figure 1 causes momentum transfer in the negative y-direction, and the force required to maintain a specific velocity V will be proportional to the area of the upper plate. Consequently, this force, when divided by the area, has the units of pressure, and is in fact the xy component, P_η, of the pressure tensor, $\mathbf{P}$. The velocity profile is characterized by the strain rate $\gamma = du_x / dy$, where u_x is the streaming velocity (net fluid velocity) in the x-direction. The viscosity η is defined as the constant of proportionality between the
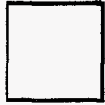
component of the pressure tensor and the strain rate,

$$P_\eta = -\eta \gamma \quad (1)$$

For so-called Newtonian fluids, the viscosity defined by Eq. (1) is independent of strain rate. It is one of the outstanding achievements of the past two decades of research in non-equilibrium molecular dynamics to have derived a rigorously correct molecular simulation algorithm to describe the molecular dynamics of a fluid element from deep within the planar Couette flow (as shown in Figure 1) while maintaining a homogenous thermodynamic state (i.e., no temperature or density gradients) in the simulation sample [2]. The algorithm, known as the SLLOD algorithm, is described in Section 2 below.
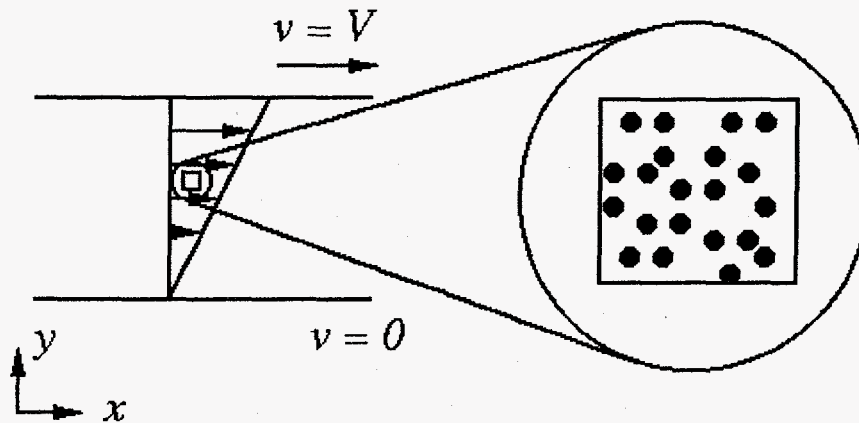

Figure 1: Geometry of planar Couette flow.

We conclude this introduction with several observations about the molecular simulation of viscosity using non-equilibrium molecular dynamics (NEMD). First, NEMD algorithms simulate into the non-linear (or non-Newtonian) regime, where the viscosity is dependent on the strain rate, and in fact decreases with increasing strain rate (shear-thinning behavior). The higher the strain rate the easier it is to calculate the viscosity, because P_η is on average zero at equilibrium and, according to Eq. (1), becomes increasingly non-zero with increasing strain rate. Thus the "signal" (average value of P_η) to thermal noise (fluctuations in P_η due to random thermal motion of the molecules) ratio is best at high strain rate, and worst at low strain rate. Secondly, for molecules which are significantly non-spherical (such as hydrocarbon chains), the dominant relaxation time for viscous motion at low strain rates is generally the rotational relaxation time of the molecule. This is because the planar Couette flow field contains a rotational part, and in order to obtain a good statistical representation of the viscous properties of a liquid at low strain rates, the simulation must encompass several rotational relaxation times. Thus, NEMD becomes difficult at low strain rates as a result of the combination of the signal-to-noise ratio

and the need to simulate long enough times to encompass a sufficient number of rotational relaxation times. Signal-to-noise ratios can be overcome either by long runs (since for a given number of molecules the statistical error is roughly proportional to $1/\sqrt{t_{sim}}$, where t_{sim} is the simulated time over which properties are sampled) or large systems (since for a given length of simulation the statistical error is roughly proportional to $1/\sqrt{N}$, where N is the number of molecules in the system). Clearly, two strategies for overcoming statistical inaccuracy in NEMD simulations are to use larger system sizes and to use longer simulation times at lower strain rates. We have used both strategies to obtain efficient, accurate NEMD simulations of the rheological properties of hydrocarbon chain fluids ($C_{10}-C_{30}$) and of a simple fluid (spherically symmetric molecules). In the former case, the rotational relaxation time is the dominant factor in obtaining good statistics, so we developed a replicated data implementation of the SLLOD algorithm to obtain good statistics at low strain rate. In the latter case, where rotational relaxation time is not an issue, domain decomposition was used to permit very large systems to be simulated (up to 364500 molecules).

We describe the NEMD simulations of hydrocarbon chains in Section 2 below, and the NEMD of very large systems of simple fluid molecules in Section 3. Our conclusions, including our future research goals, are outlined in Section 4.

2. NEMD Simulation of the Rheology of Hydrocarbon Chains

The rheological properties of liquid alkanes of intermediate molecular sizes ($C_{20}-C_{40}$) are among the most important properties in the performance of lubricants, since they constitute the lube basestock of all commercially available motor oils. Improved motor oils are produced when long-chain, frequently highly-branched hydrocarbons (C_n $n \geq 100$) are added at low dilution to improve the viscosity index (a measure of viscosity over a wide range of temperature) of the oil. The regulatorally-mandated new generation of automobiles (which will be achieved using what is collectively referred to as Clean Car Technology) will require engines operating at higher temperature to ensure more complete combustion of fuel. This will put severe demand on lubricants, and oil companies throughout the world are in the process of developing new synthetic lubricants in anticipation of these requirements. Using molecular simulation to predict the rheological properties of synthetic lubricants prior to synthesis is an extraordinarily cost-effective design tool. Using NEMD and a model for the interaction potential for liquid alkanes recently developed by the molecular modeling group at Shell Research in the Netherlands [3][4], Mundy et al. [5] and Cui et al. [6] quantitatively predicted the viscosity of normal decane at a liquid state condition similar to those in the phase envelope calculation (high temperature, moderately high density). The model was assessed by Mondello and Grest [7] for a number of equilibrium properties at near-ambient temperature conditions. We have recently implemented a multiple time step algorithm [6][8] to perform molecular simulations for chain molecules undergoing planar Couette flow on massively parallel supercomputers, particularly the Intel Paragon. This allows us to systematically study the rheological properties of liquid alkanes with molecular size of industrial interest.

The equations of motion used for liquid alkanes under planar Couette flow are the SLLOD [2] equations with Nosé dynamics [6]

$$\dot{r}_{ia} = \frac{p_{ia}}{m_{ia}} + \gamma y_{ia}\hat{x}$$
$$\dot{p}_{ia} = F_{ia} - \gamma p_{y,ia}\hat{x} - \zeta p_{ia}$$
$$\dot{\zeta} = \frac{p_\zeta}{Q}$$
$$\dot{p}_\zeta = F_\zeta \qquad\qquad (2)$$

where $\vec{r}_a$ and $\vec{p}_a$ are the position and momentum coordinates of atom a in molecule i, y_a and $P_{y,a}$ are their components, m_a is the mass, and $\hat{x}$ is a unit vector in x direction. The quantities ς, P_ς, and Q are the variables related to the Nosé thermostat, $P_\varsigma = \sum_{i\alpha} \frac{p_{i\alpha}^2}{m_{i\alpha}} - 3Nk_BT$, $Q = 3Nk_BT\tau^2$, τ is the Nosé thermostat time constant, and N is the total number of atoms in the system. Standard Lees-Edwards boundary conditions [9] are used to drive the planar Couette flow.

In a NEMD calculation, the strain-rate dependent viscosity η is determined from the constitutive relation

$$\eta = -\frac{<P_{xy}> + <P_{yx}>}{2\gamma} \quad (3)$$

where $<P_{xy}>$ and $<P_{yx}>$ are the average of the symmetric part of the xy and yx components of the pressure tensor. Further details of the method for calculating the pressure tensor are provided in Ref. [8]. Because of a combination of high frequency motions associated with intramolecular (bond vibration, bond angle bending and torsional forces) and lower frequency motion associated with intermolecular interactions, the equations of motion were solved using a multiple time step technique with Nosé constant temperature dynamics. Details can be found in the works of Tuckerman et al. [10] and Cui et al. [6]. In our implementation for this calculation, two time step sizes were used. All the internal interactions were treated as fast motions and the intermolecular interaction as the slow motion. The large time step was 2.35 fs, and the small time step was 0.235 fs. The strain rate dependent properties of the systems were calculated based on the simulation runs after the systems had reached steady state under the influence of the shear flow. This is roughly estimated to be the time for the particles at the top of the simulation cell to traverse the entire box length, about 25 ps at strain rate $\gamma = 1$ at density $\rho = 0.773 \text{ g/cm}^3$ for tetracosane and smaller for other alkanes in this study. The configuration of a neighboring higher strain rate was used as the starting configuration for the next smaller strain rate as this allows the system to reach steady state more quickly than starting from an equilibrium configuration. We have actually used relatively large times for the system to reach the steady state at each strain rate, ranging between 100 ps for the high strain rate to 470 ps for the low strain rate. The length of production runs at each strain rate were between 0.75 and 19.5 ns, which for molecules of this complexity represent extraordinary long simulations.

These simulations were performed using a replicated data code developed by us in which each processor carries a copy of the positions and velocities of all the molecules in the simulation, and the force calculation is parallelized by distribution in a load-balanced way across all the processors. At the completion of the force calculation, the results of the calculation are globally communicated to each processor. Each processor then has the force on each molecule known to it. Each processor in parallel then integrates the equations of motion of the molecules assigned to it. After the imposition of boundary conditions, and before the next force evaluation, the positions and velocities of the molecules are communicated globally to the other processors, so that each processor has a copy of the position and velocity of all the molecules. The implementation of Lees-Edwards boundary conditions is no more difficult than for any serial implementation [11]. Since the force calculation is generally by far the most time-consuming part of any molecular simulation, the replicated data formulation is a reasonably successful parallelization strategy for molecular dynamics. We regard replicated data as an effective parallelization strategy for small to medium size systems requiring long simulation times. The negative aspect of replicated data is that the wall clock time per simulation time step cannot be reduced below that required for a global communication. Thus an effective upper bound exists on the maximum number

of timesteps which can be performed.

The viscosity is plotted in Figure 2 for decane, hexadecane, and tetracosane at respective state points as a function of strain rate. These state points correspond to the experimental equilibrium density at atmospheric pressure at the respective temperatures. The figure is plotted on a log-log scale and the strain rate covers several orders of magnitude. As can be seen from the figure, the viscosity shows shear thinning behavior as a function of strain rate, typically observed in chain molecule fluid systems. At larger shear, the shear thinning follows a power law. The slope of the log-log plot varies from -0.33 to -0.41 for the four systems simulated. This can be compared with the experimentally observed slopes of -0.4 to -0.9 for polymeric fluids. This suggests that even though the alkane chains studied here are very short in comparison to polymer systems, they nevertheless exhibit some of the generality of the long-chain systems. Another interesting aspect is that the shear viscosities for decane, hexadecane, and tetracosane at high strain rate nearly overlap each other. Intuitively, this behavior should be related to the fact that at high strain rate, these fairly short and stiff alkane chains are well aligned with each other so they can slide past each other easily. In addition, the longer chain systems align with a smaller angle in the flow direction, resulting in a similar viscosity even though the chain lengths are different.
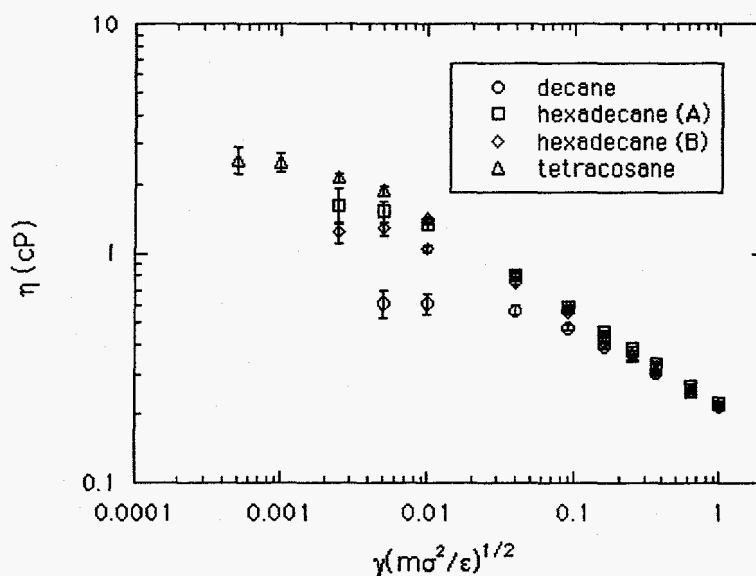

Figure 2. Viscosity as a function of strain rate for decane (temperature 298 K and density 0.7247 (g/cm³)), hexadecane (A: temperature 300 K and density 0.770 (g/cm³) , B: temperature 323 K and density 0.753 (g/cm³)), and tetracosane (temperature 333 K and density 0.773 (g/cm³)).

3. NEMD Simulation of Simple Fluids

In contrast to replicated data methods, domain decomposition techniques yield scalable molecular dynamics algorithms for simulating large systems on distributed memory parallel computers [12]. Scalability is an important factor because it allows both system size and number of processors to be increased without significantly effecting the communication costs. In a recent study, Hansen and Evans [13] proposed such a parallel algorithm for NEMD simulations of shear flow where they used domain

decomposition with a co-moving, or Lagrangian, simulation cell that changes shape (deforms) throughout the simulation.

In simulations of planar Couette flow, Lees-Edwards periodic boundary conditions have to be implemented in order to have an infinite shearing system. Parallelization of the more commonly used sliding brick form of the Lees-Edwards boundary conditions for domain decomposition algorithm results in complex communication patterns due to shifting of domains with respect to their images at the shearing boundaries [14]. The communication costs also become large with the sliding brick boundary conditions because of the rapid convection of particles through processor domains because of the shear field. Hansen and Evans eliminated these problems by using the deforming cell form of the Lees-Edwards periodic boundary conditions. In deforming cell boundary conditions, the images of the particles are taken with respect to an oblique set of coordinate axes where the y axis makes an angle θ with the vertical instead of a rigid orthogonal coordinate system used in sliding brick boundary conditions. The angle θ varies with the shear rate γ and time t as follows

$$\theta = \tan^{-1}(\gamma t) \quad (3)$$

Since the simulation cell, and hence the individual domains controlled by the processors, deforms with the shear flow under deforming cell boundary conditions, the communication patterns at the shearing boundaries are similar to those for the equilibrium molecular dynamics (EMD) case. Also the particles will diffuse through domain boundaries purely as a result of thermal motion. In order to prevent the simulation cell from deforming indefinitely in a single direction, Hansen and Evans reset the cell angle to $-45°$ whenever θ reaches $+45°$. However, in order to use link cells for the force calculations, the link cell size has to be increased from r_c, the cutoff radius in an equilibrium simulation, to $r_c/\cos(45°)$ so that the particles in a link cell interact only with particles within the same link cell or its immediate neighbors even at the maximum angle of deformation. This creates a considerable overhead in the force calculation due to additional pairs being considered. With a link cell size of $r_c/\cos(45°)$, one would consider $13.5N\rho(r_c/\cos(45°))^3$ pairs where N is the total number of particles and ρ is the number density. This is in contrast to $13.5N\rho(r_c)^3$ pairs considered in the case of an EMD where the link cell size would be r_c. In the worst case this is almost a factor of 2.8 in terms of number of operations required in force calculations. Since the force calculations are the most expensive part in any molecular dynamics simulation, this would create a large computational overhead that might offset the advantages of using the deforming cell method.

We recently proposed a modification of the deforming cell algorithm that uses a smaller maximum angle of deformation [15]. This algorithm allows us to carry out NEMD shear flow simulations with almost negligible penalty in the force calculation resulting from the deforming cell periodic boundary conditions. We also present shear viscosity results obtained using the new algorithm for the Weeks-Chandler-Andersen (WCA) system at the low shear rate regime. Large system sizes are required for these low shear rate simulations to obtain good statistical results. The SLLOD equations must be implemented with Lees-Edwards periodic boundary conditions to give an exact description of the planar Couette flow arbitrarily far from equilibrium. We use the deforming cell form of the Lees-Edwards periodic boundary conditions. Under these deforming cell boundary conditions, a particle i moves out of the simulation cell in the positive x-direction when $x_i > L + y_i \tan\theta$ where L is the length of the simulation cell and in the negative x-direction when $x_i < y_i \tan\theta$. It should be noted that θ varies with time according to Eq. (1). The conditions in other coordinate directions are the same as in case of EMD. Since the computational overhead due to the deforming cell is proportional to $(1/\cos\theta_{max})^3$ where θ_{max} is the maximum

angle of deformation, it might appear that the overhead can be minimized by shifting the cube at the shortest possible angle after talking into consideration the message passing required to remap the particles during each shifting. However, for the simulation cell to align properly with the images throughout the simulation, the shifting can be done only when the image cells immediately above the central simulation cell move through a distance that is an integral multiple of the box length. From this point of view, we can see that, in the previous study [13] the cube was realigned when the image cells moved through a distance of two box lengths (i.e. from $\theta = -45°$ to $45°$) (Figure 3A). In our algorithm, we realign the cube when the image cells move through just one box length and hence deform the simulation cell from $-26.6°$ to $26.6°$ (Figure 3B). We can see that the number of pairs considered in the worst case with this method would be 1.4 times the limiting case i.e. for the rigid simulation cell. This in contrast to 2.83 times for the previous method. Domain decomposition in our algorithm was implemented according to the link cell algorithm proposed by Pinches et. al [12].
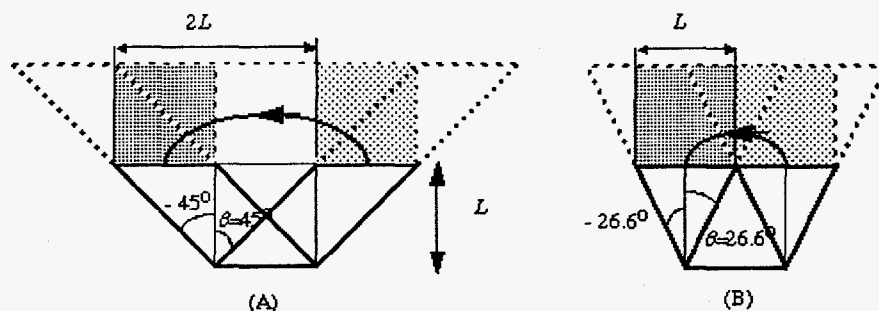

Figure 3. Angle at which the deforming cell is realigned in the Hansen and Evans algorithm (A) and in the Bhupathiraju et al algorithm (B) . The deforming simulation cell is shown with solid lines, the corresponding images are with dashed lines and the corresponding sliding brick images are shown with shaded boxes.

The simulations were carried out using the WCA fluid system whose potential is a Lennard-Jones (LJ) potential truncated at the minimum energy and shifted so that both potential and force go to zero at the cutoff. The WCA potential is given as follows.

$$\phi(r) = 4\varepsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right] + \varepsilon \quad r \leq 2^{1/6}\sigma$$
$$= 0 \quad r > 2^{1/6}\sigma$$

where ε is the well depth of the LJ potential and σ is the separation at which the interaction energy for LJ potential is zero and $r_c = 2^{1/6}\sigma$ is the cutoff distance. The fluid is at the LJ triple point, i.e. at reduced temperature, $T^* = k_B T/\varepsilon = 0.722$ and reduced density, $\rho^* = \rho\sigma^3 = 0.8442$. In this paper all the values are given in terms of the usual reduced units. The time step used in all the simulations is $\Delta t^* = (\varepsilon/m\sigma^2)^{1/2}\Delta t = 0.003$. The simulations were run on the Intel Paragon XP/S 35 and XP/S 150 supercomputers at the Oak Ridge National Laboratory.

Shear viscosity results for the WCA fluid at the L-J triple point were obtained for reduced shear rates, γ^*, in range of $0.0025 - 1.44$. For higher shear rates i.e., $\gamma^* = 0.01 - 1.44$ where the signal to noise ratio is favorable, the simulations were run for 200,000 time steps with systems of 64,000 to 108,000 particles. For low shear rates i.e., $\gamma^* = 0.0025 - 0.0081$ where the signal to noise ratio is unfavorable, the simulations were run for 400,000 time steps with systems of 256,000 to 364,000 particles. A typical run of 256,000

particles on 256 processors took between 4 and 5 hours on the Intel Paragon XP/S 35 and 150. Figure 4 shows the log-log plot of viscosity versus shear rate. Also shown in the plot is the linear viscosity result from Green-Kubo (G-K) calculations [16] and the viscosity results at two different strain rates from transient time correlation function (TTCFs) calculations [16]. The TTCFs are the nonlinear generalizations of the G-K formulas and can be used to obtain accurate viscosity results for very low shear fields with comparatively smaller system sizes. However, for each shear rate, at least 60,000 nonequilibrium starting states are needed with each running for a reduced time of 1.5. Evans and Morris reported that the TTCF results shown in Figure 3 required a total of 54 million time steps. From the figure, it is observed that the results obtained from the parallel NEMD simulations are consistent with the results from TTCF calculations. Furthermore, a clear transition from the non-Newtonian behavior to Newtonian behavior at lower shear rates is also observed. The low shear rate values also show good consistency with the zero shear rate value obtained from the Green-Kubo calculations.
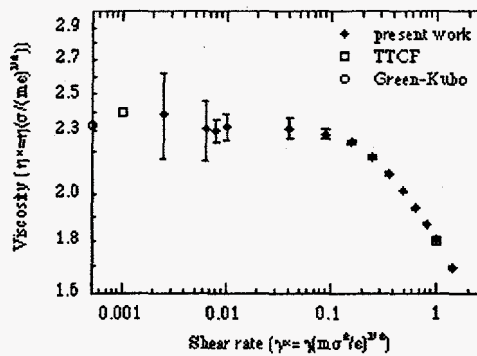

Figure 4. Shear viscosity results for the WCA fluid system.

4. Conclusions

The promise for complex molecular product design afforded by molecular simulation codes running efficiently on massively parallel is enormous. However, the limitations of this approach must also be acknowledged. The calculations reported in this paper simply could not have been performed in a timely fashion without access to the massively parallel Intel Paragons in the Center for Computational Sciences at Oak Ridge National Laboratory. For example, the lowest strain rate simulations shown in Figure 2 correspond to 550 hours of wall-clock time using 100 processors on the Intel Paragon. That is, the equivalent of almost 0.6% of all the node-hours available on the machine over a one year period. When the number of calculations to be performed is factored in, it becomes clear that the calculations reported in this paper represent a considerable investment of the total available time in the Intel Paragons at Oak Ridge National Laboratory.
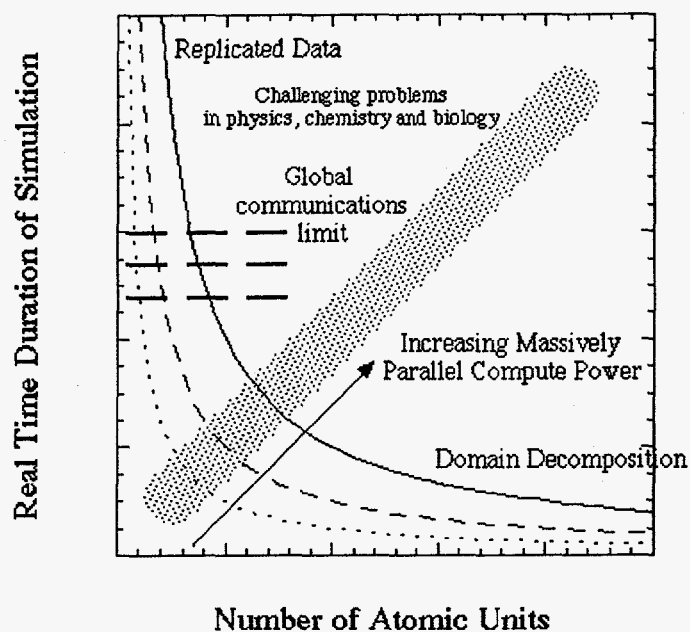
Figure 5. Trade-off between system size and total simulated time for direct molecular simulation of physical systems.

Even with this capability, many interesting problems lay outside the reach of direct molecular simulation. In Figure 5 we illustrate this with a qualitative picture that shows the trade-off between system size and total simulated time for molecular dynamics simulations on massively parallel computers. Each curve represents a new generation of massively parallel supercomputer. As we have demonstrated in this paper, domain decomposition is an effective strategy for performing simulations on large systems for relatively short times, while replicated data is the preferred method for smaller systems simulated for long times. However, replicated data is effectively limited by the time required for global communications, since no matter how much the time for the force evaluation is reduced, one cannot reduce the wall-clock time for a time step to less than the time required for two global communications. Likewise, domain decomposition remains a viable simulation strategy (i.e., exhibits scaling) only if the number of atomic units being simulated on each processor is large enough to diminish the message-passing component of the calculation to a small fraction of the elapsed wall-clock time, thus effectively limiting the possible total simulated time rather significantly. A modest improvement can be achieved by a combination of domain decomposition and replicated data, and we are actively implementing such codes in our research group.

The crucial difficulty in expanding substantively the range of problems to which molecular simulation is applicable is that the interesting problems in chemistry and biology (such as polymer rheology, micelle self-assembly and protein folding) all lie along the axis shown in the figure which corresponds both to increasing numbers of atomic units and increased simulation time required. This is because increasing the number of atomic units in a real system invariably increases the relaxation times for the properties of

interest. To fully achieve our goal of routine molecular product design using massively parallel supercomputers will demand the use of statistical mechanical theory which can guide automated coarse-graining of the molecular detail during the course of a simulation, thus permitting much larger timesteps. This is the current focus of our research efforts in molecular simulation of complex systems on massively parallel supercomputers.

Acknowledgments

References

[1]Allen, M.P. and D.J. Tildesley, *Computer Simulation of Liquids*. 1987, Oxford: Oxford University Press.

[2]Evans, D.J. and G.P. Morriss, *Statistical Mechanics of Nonequilibrium Liquids*. 1990, New York: Academic Press.

[3]Siepmann, J.I., S. Karaborni, and B. Smit, *Vapor Liquid Equilibria of Model Alkanes*. Journal of the American Chemical Society, 1993. **115**: p. 6454.

[4]Siepmann, J.I., S. Karaborni, and B. Smit, *Simulating the Critical Behaviour of Complex Fluids*. Nature, 1993. **365**: p. 330.

[5]Mundy, C.J., J.I. Siepmann, and M.L. Klein, *Calculation of the shear viscosity of decane using a reversible multiple time-step algorithm*. Journal of Chemical Physics, 1995. **102**: p. 3376.

[6]Cui, S.T., P.T. Cummings, and H.D. Cochran, *Multiple time step nonequilibrium molecular dynamics simulation of n-decane*. Journal of Chemical Physics, 1996. **104**: p. 255.

[7]Mondello, M. and G.S. Grest, *Molecular dynamics of linear and branched alkanes*. Journal of Chemical Physics, 1995. **103**: p. 7156.

[8]Cui, S.T., *et al.*, *Molecular dynamics simulations of the rheological properties of normal decane, hexadecane and tetracosane*. J. Chem. Phys., 1996. accepted for publication.

[9]Lees, A.W. and S.F. Edwards, *The computer study of transport processes under extreme conditions*. J. Phys. C: Solid State, 1972. **5**: p. 1921.

[10]Tuckerman, M.E., B.J. Berne, and G.J. Martyna, *Reversible multiple time step molecular dynamics.* J. Chem. Phys., 1992. **97**: p. 1990.

[11]Evans, D.J. and G.P. Morriss, *Non-Newtonian molecular dynamics.* Computer Phys. Reports, 1984. **1**: p. 297.

[12]Pinches, M.R.S., D.J. Tildesley, and W. Smith, *Large scale molecular dynamics on parallel computers using the link-cell algorithm.* Molecular Simulation, 1991. **6**: p. 51.

[13]Hansen, D.P. and D.J. Evans, *A parallel algorithm for nonequilibrium molecular dynamics simulation of shear flow on distributed memory machines.* Molecular Simulation, 1994. **13**: p. 375-393.

[14]Rastogi, S. and N. Wagner, *Massively Parallel Non-Equilibrium Brownian Dynamics Simulations for Complex Fluids: The Rheology of Brownian Suspensions.* Computers chem. Engng, 1995. **19**: p. 693-718.

[15]Bhupathiraju, R.K., P.T. Cummings, and H.D. Cochran, *An Efficient Parallel Algorithm for Nonequilibrium Molecular Dynamics Simulations of Very Large Systems in Planar Couette Flow.* Molec. Phys., 1996. in press.

[16]Evans, D.J. and G.P. Morriss, Phys. Rev. A, 1988. **38**: p. 41.

Biographies

Ravi Bhupathiraju

Ravi Bhupathiraju is a graduate student in the Department of Chemical Engineering at the University of Tennessee, Knoxville performing research in the area of large scale molecular simulations on parallel supercomputers with main emphasis on nonequilibrium molecular dynamics simulations. He received his B.S. in Chemical Engineering from the Indian Institute of Technology, Madras in 1993 and M.S. also in Chemical Engineering from the University of Tennessee, Knoxville in 1996. He has published 2 articles in refereed journals. He is a member of American Institute of Chemical Engineers.

e-mail: ravib@verdi.engr.utk.edu
WWW: http://flory.engr.utk.edu/~ravi

S. T. Cui

Shengting Cui is a research associate in the Department of Chemical Engineering at the University of Tennesse and Chemcial Technology Division at Oak Ridge National Laboratory, where he works on the molecular simulation of rheological properties of industrially interesting lubricant molecules and mixing properties of polymer blends. He received his B.S. in Geophysics from the University of Science and Technology of China in 1981, his PH.D in Engineering Physics from University of Virginia in 1989. He worked as a postdoctoral researcher at Australian National University from 1989-1992 and at Massachusetts Institute of Technology from 1992-1994. He has published over 20 papers in refereed

journals. He is a member of the American Institute of Chemical Engineers.

e-mail: cui@utkux1.utk.edu
WWW: http://flory.engr.utk.edu/~cui

Sunil Gupta

Sunil Gupta is a PhD candidate in the Department of Chemical Engineering at the University of Tennessee, Knoxville. His research in the area of nonequilibrium molecular dynamics focuses on chain molecule systems confined between walls. He received his B.S. in Chemical Engineering from the University of Bombay, India in 1989 and M.S. in Chemical Engineering from the State University of New York, Buffalo in 1992. He is a member of the American Institute of Chemical Engineers.

e-mail: sunil@verdi.engr.utk.edu
WWW: http://flory.engr.utk.edu/~sunil

Hank D. Cochran

Hank Cochran is a group leader and senior research staff member in the Chemical Technology Division at Oak Ridge National Laboratory where he performs research on the applications of statistical mechanics in chemical engineering, especially research involving supercritical fluids and involving molecular simulations. In addition, since 1985, he has been Adjunct Professor of Chemical Engineering at the University of Tennessee, Knoxville. He received the BS degree in chemical engineering from Princeton University in 1965 and the MS and PhD degrees, also in chemical engineering, from MIT in 1967 and 1973. He has published more than 65 articles in refereed journals (plus more than 100 technical reports), and he holds two patents. He regularly organizes and chairs sessions at national and international meetings. He has served on scientific and technical review boards for the International Energy Agency, the U. S. Department of Energy, the Gas Research Institute, and other organizations. He serves as an evaluator for the Accreditation Board for Engineering and Technology. He is a member of the American Institute of Chemical Engineers, the American Chemical Society, the American Physical Society, and Sigma Xi.

e-mail: hdc@ornl.gov
WWW: http://www.ornl.gov/divisions/ctd/EnergyResearch/SolutionThermo/Cochran.html

Peter T. Cummings

Since 1994, Peter Cummings has held a joint appointment as Distinguished Professor of Chemical Engineering at the University of Tennessee-Knoxville and as Distinguished Scientist in the Chemical Technology Division at Oak Ridge National Laboratory. His research is focused on the application of statistical mechanics and molecular simulation to physical properties prediction and biophysical processes, chemical process optimization, molecular simulation on massively parallel supercomputers, and mathematical modeling of bacterial migration processes related to *in situ* bioremediation. He received the B. Math. degree in Mathematics (with First Class Honors and University Medal) from the University of Newcastle (Australia) in 1976 and the Ph.D. degree in applied mathematics and statistical

mechanics from the University of Melbourne (Australia) in 1980. In addition to post-doctoral positions in Physics at the University of Guelph in Canada and in Chemistry at the State University of New York at Stony Brook, he taught and conducted research in chemical engineering at the University of Virginia from 1983 to 1993 before assuming his current position. He has published more than 150 articles in refereed journals. He regularly organizes and chairs sessions at national and international meetings, and is a frequent keynote speaker at such meetings. He has served on editoral boards of Fluid Phase Equilibria and Molecular Physics. He is a member of the American Institute of Chemical Engineers, the American Chemical Society, the American Physical Society, Society of Rheology, Materials Research Society, and Sigma Xi.

e-mail: ptc@utk.edu
WWW: http://flory.engr.utk.edu/ptc/bio.html

DISCLAIMER