INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS
Int. J. Numer. Meth. Fluids 2005; **48**:1337–1349
Published online 12 May 2005 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/fld.983

Non-equilibrium behaviour of equilibrium reservoirs in molecular simulations

Martin W. Tysanner*,† and Alejandro L. Garcia‡

Department of Physics, San Jose State University, San Jose, CA 95192-0106, U.S.A.

SUMMARY

We explore two widely used algorithms for fluid reservoirs in molecular simulations and demonstrate that they may induce non-physical non-equilibrium effects, even in systems that should be at equilibrium. For example, correlations of momentum and density fluctuations lead to a bias in the mean fluid velocity when measured as the mean over samples of instantaneous fluid velocity. The non-physical behaviour is entirely computational in origin and is an instance of a more general issue in molecular simulations: a failure to correctly model stochastic properties may induce non-equilibrium behaviour that does not exist in the corresponding physical system. Finally, we demonstrate that simple algorithm corrections eliminate this artifact. Copyright © 2005 John Wiley & Sons, Ltd.

KEY WORDS: molecular simulations; reservoirs; fluctuations; direct simulation Monte Carlo

1. INTRODUCTION

Molecular simulations have proved useful in the study of a wide variety of problems in fluid mechanics [1]. In simulations of an open system or a composite of two or more concurrent simulations, it is common to employ particle reservoirs. One example is a molecular simulation that mimics a 'virtual wind tunnel' with reservoirs maintaining the inflow and outflow conditions [2, 3]; another is the use of a reservoir in a hybrid simulation to couple a partial differential equation solver and a molecular computation [4, 5]. Although reservoirs are conceptually straightforward and their basic algorithms are well-established [6], we have found subtle, computationally induced non-equilibrium behaviour in equilibrium systems employing two common reservoir algorithms. It is the purpose of this paper to point out the existence of this non-physical effect, describe its origin, and present ways to avoid it.

Figure 1 illustrates a particularly simple system consisting of a dilute gas in a box with a reservoir at the left boundary and a specular wall at the right. The state of the reservoir is

*Correspondence to: Martin W. Tysanner, Department of Physics, San Jose State University, San Jose, CA 95192-0106, U.S.A.
†E-mail: tysanner@pacbell.net
‡E-mail: algarcia@algarcia.org

Received 24 August 2004
Revised 4 March 2005
Accepted 7 March 2005

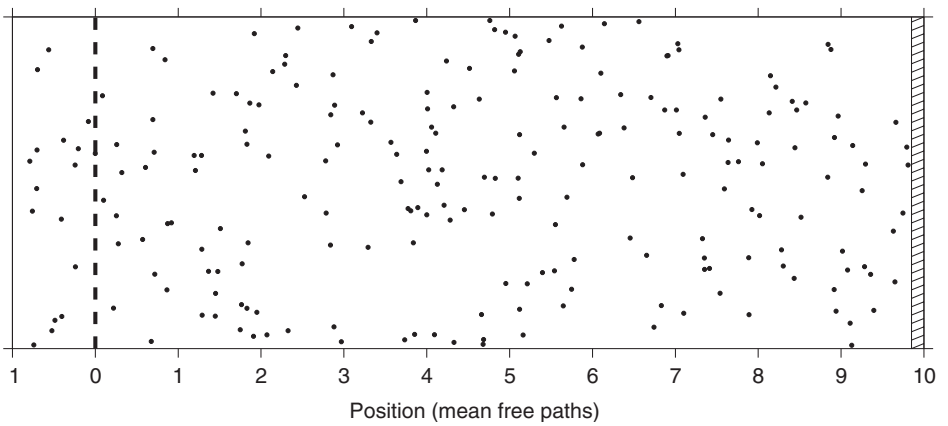

Figure 1. Conceptual diagram of the system with a reservoir at the left boundary and a specular wall at the right boundary. Boundaries are periodic in the y- and z-directions. Points represent individual fluid molecules; simulations typically used thousands of particles.

maintained at thermodynamic equilibrium with a given density and temperature. Specifically, at each time step a fixed number of molecules are randomly placed into the reservoir volume; their velocities are randomly assigned from a Maxwell–Boltzmann distribution with zero mean velocity. Particles move in and out of the reservoir, but any particles within the reservoir at the end of a time step are discarded. The specular wall at the right boundary precludes any net fluid flow, so the system should yield a stationary fluid at equilibrium.

Figure 2 shows the results of the corresponding simulation of this simple system. Note that a non-zero mean fluid velocity, $\langle u^x \rangle_s$, is clearly evident near the reservoir. This anomalous velocity is obtained (after a lengthy relaxation time to remove possible initial transients) as an average of samples of instantaneous fluid velocity. The effect is small but statistically significant; error bars are estimated using Reference [7].

Although we used the direct simulation Monte Carlo (DSMC) method [8, 9] for all simulations presented here, the nature of the effect in Figure 2 is generic to molecular simulations. The anomalous fluid velocity should also appear with other particle simulation methods like molecular dynamics (MD) [10] or lattice gases (LG) [11]. Our use of DSMC allowed us the luxury of decoupling the reservoir implementation from the main system, a simplification that is not entirely possible with MD because of its need to evaluate molecular interactions across the interface between the reservoir and the system. For MD simulations, the phenomenon we demonstrate is one more factor to consider in an already-delicate process of reservoir design [12, 13].

We used sample averaging measurement (SAM) to measure the anomalous fluid velocity in Figure 2. In the previous work [14] we determined that SAM is susceptible to bias under non-equilibrium conditions. In that paper we also showed that the origin of the bias is a physically occurring spatial correlation of non-equilibrium fluctuations in fluid density and momentum, and that an alternative definition of the mean, i.e. cumulative average measurement (CAM), is not susceptible to the bias. In contrast to SAM, which sums the samples of instantaneous fluid *velocity*, CAM obtains the mean velocity, $\langle u^x \rangle_c$, by summing samples of instantaneous

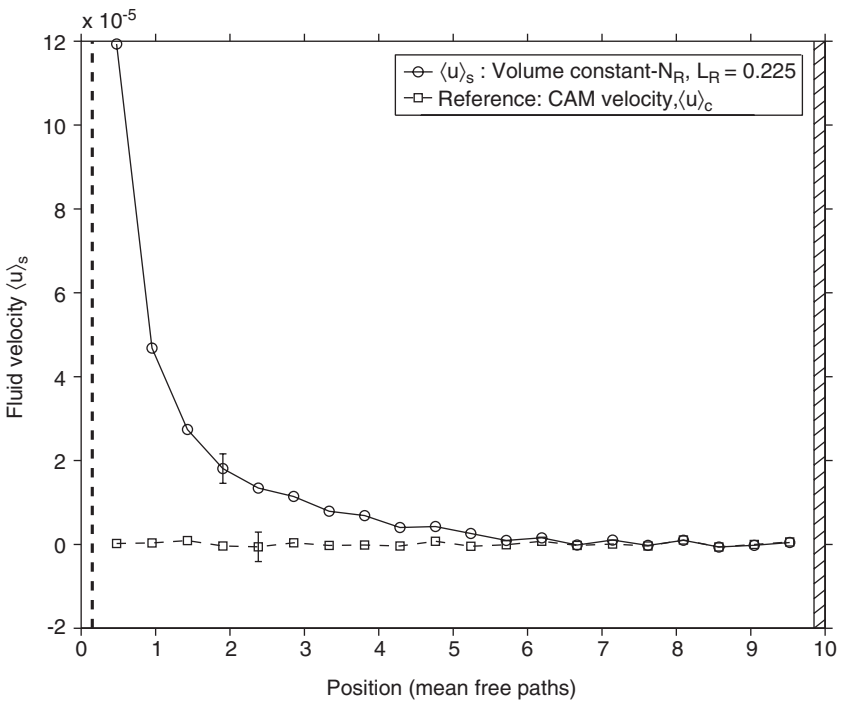

Figure 2. Equilibrium system at $T = 5$ with a reservoir at the left boundary and a specular wall at the right boundary. The x-axis in this and subsequent plots corresponds to Figure 1. Mean fluid velocity is computed from 4×10^8 samples by: (circles) sample averaging measurement (SAM), $\langle u^x \rangle_s$; and (squares) cumulative average measurement (CAM), $\langle u^x \rangle_c$. See Equations (2) and (1), respectively, for definitions. Note that SAM measures an anomalous flow from the reservoir into the system, even though the unbiased CAM shows there is no actual flow.

fluid *momentum* divided by the cumulative mass over all samples. At equilibrium the SAM and CAM definitions of mean fluid velocity are equivalent since there are, of course, no non-equilibrium correlations of fluctuations.

Even though the boundary conditions imply an equilibrium system, Figure 2 shows distinct differences between the SAM and CAM fluid velocity profiles. The mass flux is zero since it depends on $\langle u^x \rangle_c$, yet the non-zero SAM velocity is indicative of a non-equilibrium system. We will show that this arises from a subtle deficiency in the reservoir algorithm: compared to the true physical system, the reservoir algorithm induces an excessive correlation between fluctuations in fluid density and momentum, which is revealed by the SAM velocity.

After giving necessary definitions and explaining important features of our simulation system, we will consider two different, practical reservoir algorithms and evaluate their performance in correctly modelling a physical reservoir. We employ the SAM definition of fluid velocity as a tool to reveal non-physical behaviour. The two algorithms or minor variants are in common use and both can produce undesirable correlations. We will demonstrate that a

minor change to each of them can eliminate these correlations without noticeably increasing computation time.

2. BACKGROUND AND NOTATION

A perfect reservoir has infinite capacity. That is, the addition or removal of a particle does not affect the fluid density, momentum or temperature within the reservoir. In practice, efficient reservoir algorithms in molecular simulations employ a relatively small number of particles, so the fluid within the reservoir will be susceptible to non-physical variations of these quantities if two general requirements are not met. Broadly stated, the fluid in the reservoir must not in any way respond to changes within the main system. Additionally, the density and velocity distribution profile of particles leaving the reservoir must have the same mean and fluctuation as would occur in a corresponding physical system.

For a dilute gas it is simple to meet the first requirement: whenever a particle in the system reaches a boundary that is defined by a reservoir, it is simply eliminated (this is more subtle in MD simulations; more on that later). This precludes any mechanism for the system to affect the reservoir. The second requirement, ensuring that the flux of particles leaving the reservoir corresponds to what would occur in the physical system, is more subtle. Maintaining the necessary mean density, temperature and fluid momentum is the first and most important step, but issues also arise with respect to the instantaneous fluctuations of these same quantities, as we will demonstrate.

We carry over the terminology and notation conventions, summarized below, from our preceding work [14]. We assume a molecular simulation that partitions the system into K cells of equal volume and gathers statistics separately for each cell. We identify a particular cell with a subscript. Cell k has volume V_k; during a time t_j it contains $N_k(t_j)$ particles. For a single species fluid all particles have mass m, so $M_k = mN_k$ is the instantaneous fluid mass in the cell. The system evolves in discrete time steps of duration τ, where τ is the interval between successive times t_j and t_{j+1}. Following an initial relaxation period to remove transient effects from the initial conditions, a total of S samples of hydrodynamic and thermodynamic variables like fluid momentum, temperature and particle number are taken, typically at successive time steps.

Our present investigations focused on a one-dimensional system that was conceptually similar to that shown in Figure 1. The right boundary along the primary axis (the x-axis) was a specular wall; the left boundary was a reservoir, the construction of which depended on the simulation. Boundaries were periodic in the y- and z-directions. The total number of particles in the system, $N_\Sigma(t_j) = \sum_{k=1}^{K} N_k(t_j)$, varied with time because the reservoir created an open system.

A typical simulation had the following characteristics. The gas was a single species of molecular mass $m = 1$, consisting of hard spheres of diameter $d = 7.5 \times 10^{-2}$. The mean number of particles in the system during a simulation was $\langle N_\Sigma \rangle = 2000$. The mean temperature was constant over time and position; as in Reference [14] the reference temperature was $T = 1$, but for the present study $T = 5$ was generally preferred because the higher molecular speeds exaggerated the magnitude of the effect under study. The system contained $K = 20$ cells distributed along a total length $L = 2.25$, with a cell volume $V_k \cong 0.57$. With these parameters the number density was $n = N_k/V_k \cong 175$, corresponding to a volume fraction of 3.4×10^{-2};

for the mean free path, $\lambda \cong 0.23$. Boltzmann's constant was $k_B = 0.5$, so the most probable molecular speed, $v_{mp} = \sqrt{2k_B T/m}$, was unity at $T = 1$.

The typical simulation gathered statistics for each of at least $S = 5 \times 10^6$ successive DSMC time steps of duration $\tau = (L/K)/5v_{mp}$. The initial system relaxation period was 5×10^6 time steps, corresponding to about 5.6×10^8 collisions.

As mentioned in the Introduction, there are two common definitions of mean fluid velocity; we denote these as CAM and SAM. The CAM definition for a single species fluid is

$$\langle \mathbf{u}_k \rangle_c = \frac{(1/S)\sum_{j=1}^{S}\sum_{i\in k}^{N_k(t_j)} m\mathbf{v}_i(t_j)}{(1/S)\sum_{j=1}^{S} N_k(t_j)} = \frac{\langle \mathbf{J}_k \rangle}{\langle M_k \rangle} \tag{1}$$

where

$$\langle M_k \rangle = m\langle N_k \rangle = \frac{m}{S}\sum_{j=1}^{S} N_k(t_j)$$

and

$$\langle \mathbf{J}_k \rangle = \frac{1}{S}\sum_{j=1}^{S} \mathbf{J}_k(t_j) = \frac{1}{S}\sum_{j=1}^{S}\sum_{i\in k}^{N_k(t_j)} m\mathbf{v}_i(t_j)$$

are the mean fluid mass and momentum, respectively, in cell k. The notation $i \in k$ denotes the index i ranging over the particles within cell k. The SAM definition is the mean over all samples of the average particle velocity in a cell at each sample, $\bar{\mathbf{v}}_k(t_j)$:

$$\langle \mathbf{u}_k \rangle_s = \frac{1}{S}\sum_{j=1}^{S} \bar{\mathbf{v}}_k(t_j) = \frac{1}{S}\sum_{j=1}^{S} \frac{1}{N_k(t_j)} \sum_{i\in k}^{N_k(t_j)} \mathbf{v}_i(t_j) = \left\langle \frac{\mathbf{J}_k}{M_k} \right\rangle \tag{2}$$

$$= \left\langle \frac{\langle \mathbf{J}_k \rangle + \delta\mathbf{J}_k}{\langle N_k \rangle + \delta N_k} \right\rangle = \langle \mathbf{u}_k \rangle_c \left(1 + \frac{\langle (\delta N_k)^2 \rangle}{\langle N_k \rangle^2} \right) - \frac{\langle \delta\mathbf{J}_k\, \delta N_k \rangle}{m\langle N_k \rangle^2} + O([\delta X]^3) \tag{3}$$

where $O([\delta X]^3)$ are terms of cubic order in the fluctuations. Note that the average particle velocity $\bar{\mathbf{v}}_k(t_j)$ within the cell at time t_j is equivalent to the instantaneous fluid velocity $\mathbf{u}_k(t_j)$. Since $\mathbf{J}_k = mN_k\mathbf{u}_k$ and thus $\delta\mathbf{J}_k = m\langle N_k \rangle \delta\mathbf{u}_k + m\langle \mathbf{u}_k \rangle_c \delta N_k$, the expression for $\delta\mathbf{J}_k$ can be substituted into the right side of Equation (3) to give

$$\langle \mathbf{u}_k \rangle_s = \langle \mathbf{u}_k \rangle_c - \frac{\langle \delta\mathbf{u}_k \delta N_k \rangle}{\langle N_k \rangle} + O([\delta X]^3) \tag{4}$$

The covariance $\langle \delta\mathbf{u}_k \delta N_k \rangle$ is zero at equilibrium [15], but at non-equilibrium it evidently biases $\langle \mathbf{u}_k \rangle_s$; the CAM definition remains unbiased under all conditions. Thus, the statistically significant difference between $\langle \mathbf{u}_k \rangle_s$ and $\langle \mathbf{u}_k \rangle_c$ in Figure 2 strongly indicates that the constant temperature, steady state system was not truly at equilibrium.

In the next sections we consider two different, practical reservoir algorithms that are probably the most commonly used. Although all discussion is for one-dimensional systems (i.e. periodic in y and z), the generalization to two or three-dimensional systems is straightforward. We employ the SAM definition of fluid velocity to help evaluate the performance of

these various algorithms in modelling a physical reservoir, in particular to detect an erroneous, computationally induced correlation of fluctuations.

3. VOLUME GENERATION RESERVOIRS

The volume generation reservoir algorithm fills a fixed volume at the beginning of each time step by random placement of particles with randomly assigned, appropriately distributed velocities (see Figure 1). We only consider the case where the mean fluid velocity in the reservoir is zero; extension to the case of a bulk flow is straightforward and outlined in Reference [6], which calls volume generation reservoirs the *alternative method* for stream boundary conditions. Those particles that reach the system within the time step become part of it; the remainder are discarded. Refreshing the reservoir with an entirely new set of particles at each time is necessary to mimic a reservoir of infinite capacity. The only subtlety is ensuring that the reservoir depth $L_R = v^x_{max}\tau$ is sufficient so that it is extremely improbable that any particle could travel more than the length of the reservoir within a single time step; otherwise the finiteness of the reservoir will bias the velocity distribution of particles entering the system by cutting off the high velocity tail of the distribution [6]. We found that setting v^x_{max} to six standard deviations above the most probable speed v_{mp} was insufficient, but that 13 standard deviations ($v^x_{max}\tau$ of two cell widths in our system) was more than adequate.

To simulate an equilibrium reservoir of volume V_R at number density n_R and temperature T_R, the volume generation algorithm proceeds as follows:

0. Determine the number of particles to be generated within the reservoir as $N_R = n_R V_R$. This is the *constant-N_R* variant of the algorithm.
1. Assign each particle a random, uniformly distributed position within the reservoir volume.
2. Randomly assign the velocity components of each particle by the Maxwell–Boltzmann distribution [8]

$$P(v) = \sqrt{\frac{m}{2\pi k_B T_R}}\, e^{-mv^2/2k_B T_R} \tag{5}$$

3. Move the particles and discard any that do not enter the system during the time step.

Three technical points bear mention. First, the number of reservoir particles, N_R, is necessarily discrete though the expected mean, $n_R V_R$ is not, in general, an integer. Simple rounding may create an undesirable bias; the correct mean is reliably obtained by 'random rounding,' $N_R = \mathbf{floor}(n_R V_R + \mathscr{R})$ were $\mathscr{R}$ is a uniformly distributed random deviate. Second, a possible optimization is to recognize that, on average, half the particles in the reservoir volume are moving away from the reservoir interface and will automatically be discarded. For the geometry shown in Figure 1, efficiency improves if N_R is halved (before rounding) and v^x is taken as positive. Third, for dilute gas simulations, for example our DSMC computations, the particles move independently and thus may be generated and moved one by one instead of generating all N_R at once.

We used the above procedure in the simulation of Figure 2. From continuity arguments, the presence of a specular wall at the end of the system opposite the reservoir precluded any mean flow within the system, and the CAM fluid velocity profile in Figure 2 confirms its absence.

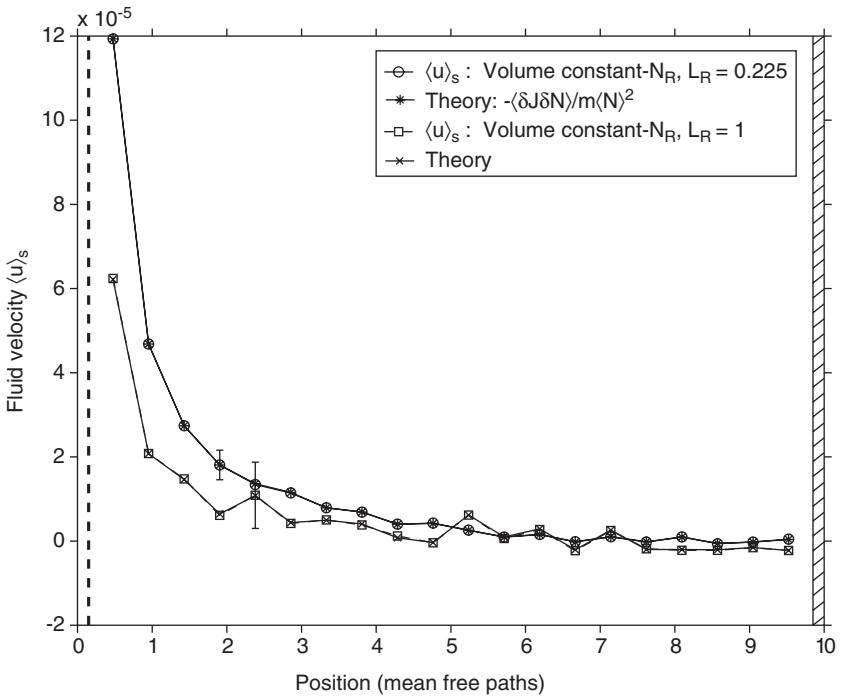

Figure 3. SAM fluid velocity, $\langle u^x \rangle_s$, from Figure 2 (circles) and the measured $\langle \delta \mathbf{J}\, \delta N \rangle$ covariance term in Equation (3) (asterisks), versus position. Their close match indicates that the anomalous mean fluid velocity is entirely due to correlations of non-equilibrium fluctuations. Evidently, lengthening the reservoir from $L_R = 0.225$ to 1.0 reduces the magnitude of the effect ($\langle u^x \rangle_s$: squares; Equation (3): x-marks).

Furthermore, the reservoir was at constant temperature and it maintained the system at the steady state. These considerations indicate that the system should be at equilibrium, and yet the simulation results showing a non-zero SAM fluid velocity are those of a non-equilibrium system.

Interestingly, the anomalous flow $\langle \mathbf{u} \rangle_s$ measured by SAM is indicative of the kinds of correlations of non-equilibrium fluctuations observed in systems with a strong temperature gradient [14]. To test this hypothesis, we measured the $\langle \delta \mathbf{J}\, \delta N \rangle$ covariance term in Equation (3); because $\langle \mathbf{u} \rangle_c = 0$ in our simulations the $\langle (\delta N)^2 \rangle$ term did not contribute. The plot of the covariance term in Figure 3 matches the $\langle \mathbf{u} \rangle_s$ profile very closely, implying non-equilibrium conditions.

An independent measurement provides further evidence that the *constant-N_R* variant of the volume generation algorithm leads to non-equilibrium behaviour. From equilibrium statistical mechanics [15] we expect the number of particles $N_{\rightarrow}$, leaving a reservoir at equilibrium during a fixed time interval, to be Poisson-distributed, and similarly for the number of particles leaving the system, $N_{\leftarrow}$. Figure 4 shows the relevant histograms of $N_{\rightarrow}$ and $N_{\leftarrow}$ for the

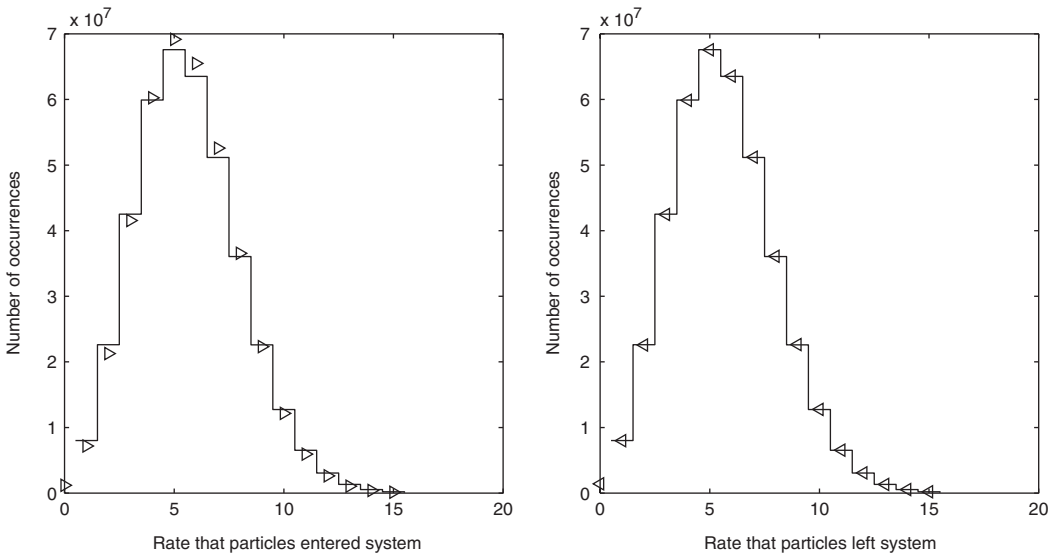

Figure 4. Histograms of $N_{\rightarrow}$ (left figure) and $N_{\leftarrow}$ (right), the number of particles leaving the reservoir and leaving the system, respectively, during a time step. For the simulation $\sum_{j=1}^{S}[N_{\rightarrow}(t_j)-N_{\leftarrow}(t_j)]=43$; the total number of particles, $\sum_{j=1}^{S}N_{\rightarrow}(t_j)$ and $\sum_{j=1}^{S}N_{\leftarrow}(t_j)$, differ by about two parts in 10^8. The number of particles N_{R} generated in the reservoir at each time was constant, with $T_{\mathrm{R}}=5$ and density n_{R} yielding $\langle N_{\Sigma}\rangle=2000$. Apparently, the distribution of $N_{\rightarrow}$ was not Poisson, in contrast to $N_{\leftarrow}$, as a comparison with the Poisson distribution stair-step plot shows.

simulation of Figure 2. The corresponding Poisson distribution is also plotted for comparison with theory. The distribution of $N_{\leftarrow}$ closely follows Poisson statistics, but there are noticeable deviations in the distribution for particles leaving the reservoir. Although the reservoir maintained a constant density and temperature and its particles had a Maxwell–Boltzmann speed distribution, it evidently does *not* model an equilibrium reservoir. Since there was no way for the system to affect the reservoir, the source of the non-equilibrium behaviour must lie in the reservoir algorithm itself.

The deviation from Poisson statistics suggests one way to improve the algorithm. By replacing the calculation of N_{R} in step (0) with a Poisson-distributed random number about the mean value $\langle N_{\mathrm{R}}\rangle=n_{\mathrm{R}}V_{\mathrm{R}}$, the algorithm generates the distribution of N_{R} given by the grand canonical ensemble [15]. This modification yields the *Poisson-*N_{R} variant of the volume generation algorithm. The SAM fluid velocity profile in Figure 5 shows that the modification eliminated the SAM velocity bias entirely. An examination of the corresponding histogram (not shown) demonstrated that particles now entered the system according to the same Poisson statistics as those leaving the system. Given that the reservoir depths, temperatures and other reservoir physical parameters were the same for the simulations of both Figures 3 and 5 it indicates that the anomalous SAM velocity was *not* caused by the reservoir being too shallow (i.e. insufficiently large v_{max}^{x}).

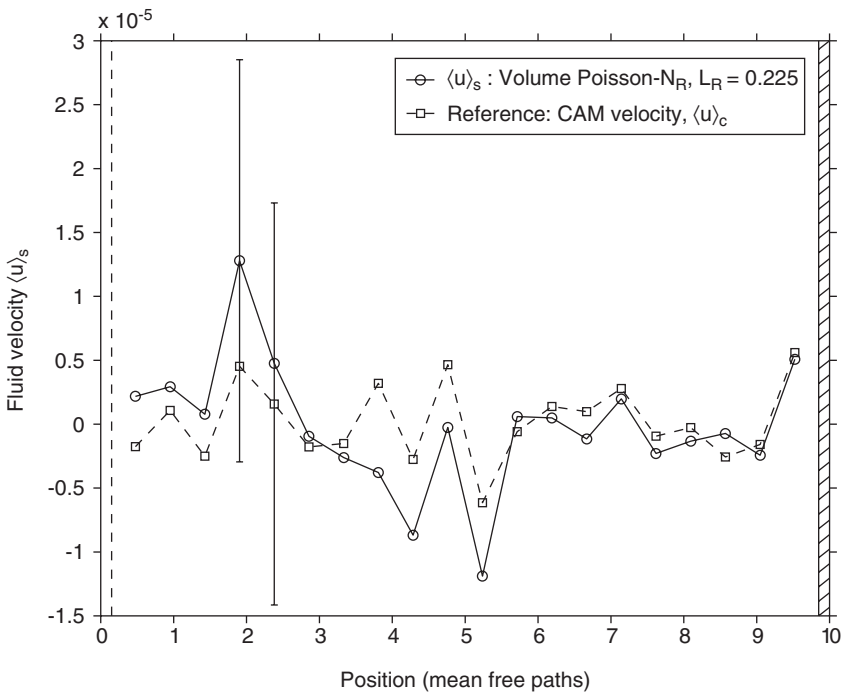

Figure 5. Sample averaged measurement of mean fluid velocity, $\langle u_k^x \rangle_s$, using the *Poisson-*N_R variant of the volume generation algorithm for the system of Figure 2, obtained over $S = 10^8$ samples. Comparison with that figure shows that this modification to the reservoir algorithm entirely eliminated the anomalous fluid flow, fully consistent with expectations for a system at equilibrium.

A final observation is relevant to the non-equilibrium behaviour of the *constant-*N_R variant. Although the *Poisson-*N_R variant corrected the algorithm by introducing a simple change in how N_R is computed in step (0), any change that reduces the time correlation of the number of particles entering the system should improve the original *constant-*N_R variant. For example, using simulations we found that reducing the time step proportionately reduced the anomalous SAM fluid velocity. Making the reservoir deeper had a similar proportional effect; in fact when the reservoir depth equalled the system length, the SAM and CAM velocity profiles were statistically indistinguishable with $S = 5 \times 10^6$ samples, although there was a slight deviation from Poisson statistics in the histogram. Even eliminating the initial division by two of N_R, so that particles were generated with either positive or negative v^x, reduced the magnitude of $\langle \mathbf{u} \rangle_s$ by a factor of two. Given that each of these changes increases the computation time without entirely removing the non-equilibrium behaviour except in the limit $L_R \to \infty$, the modification to obtain the *Poisson-*N_R algorithm seems optimal. In fact, the simulation of Figure 5 employed the optimization of halving N_R and generating all particles with $v^x \geqslant 0$, but no resulting non-equilibrium behaviour was evident.

4. SURFACE GENERATION RESERVOIRS

A surface generation reservoir is essentially a surface that emits particles at a rate that is determined by the required fluid density and temperature. As with thermal walls, the velocity distribution is the one-sided (biased) Maxwell–Boltzmann distribution for the normal component [8]:

$$P(v^x) = \frac{m}{k_B T_R} v^x e^{-m(v^x)^2/2k_B T_R} \tag{6}$$

and the ordinary Maxwell–Boltzmann distribution Equation (5) for the parallel components. As in the previous section, we only consider the case of an equilibrium reservoir of stationary fluid; the extension to the case of a bulk flow is straightforward and outlined in Reference [6], which calls surface generation reservoirs the *standard method* for stream boundary conditions.

The surface generation algorithm differs from volume generation in two aspects. First, all $N_{\rightarrow}$ particles are generated at the reservoir's surface so the distance $v^x\tau$ that an emitted particle travels during a time step must be scaled by a random deviate to correctly simulate its origination at a random time from an infinitely deep reservoir. Second, *all* generated particles enter the system; this makes surface generation reservoirs computationally more efficient than volume generation when the mean velocity of the reservoir is zero.

To simulate an equilibrium reservoir at number density n_R and temperature T_R, the surface generation algorithm proceeds as follows:

0. Determine the number of particles to be generated for the time interval $\Delta t = \tau$, given the expected mean [8]

$$\langle N_{\rightarrow} \rangle = n_R A_R \sqrt{\frac{k_B T_R}{2\pi m}}\, \Delta t \tag{7}$$

   where A_R is the surface area of the reservoir.
1. Assign each particle a random, uniformly distributed position on the surface (unnecessary in one-dimensional systems).
2. Randomly assign the velocity components of each particle by the biased Maxwell–Boltzmann distribution.
3. Move each particle at its assigned velocity. The final x-coordinate is scaled by a uniform random deviate, i.e. $x = \mathscr{R} v^x \tau$.

The algorithm has two variants that differ only in step (0). Since a discrete value of $N_{\rightarrow}$ is required one may take $N_{\rightarrow} = \mathbf{floor}(\langle N_{\rightarrow} \rangle + \mathscr{R})$, i.e. 'random rounding;' this is the *constant-*$N_{\rightarrow}$ variant. The *Poisson-*$N_{\rightarrow}$ variant generates a random value for $N_{\rightarrow}$ as a Poisson-distributed number with the desired mean, $\langle N_{\rightarrow} \rangle$.

Figure 6 compares the two surface generation variants. While the *Poisson-*$N_{\rightarrow}$ case shows no undesirable non-equilibrium correlations, *constant-*$N_{\rightarrow}$ clearly performs very poorly in that regard, an order of magnitude worse than the *constant-*N_R variant of the volume generation algorithm. Relating this to their respective histograms, the *constant-*$N_{\rightarrow}$ case has at most two non-zero points, corresponding to the two integers on either side of the mean, while the histogram for the *Poisson-*$N_{\rightarrow}$ variant necessarily matches a Poisson distribution exactly.

As pointed out in Reference [6], the volume generation reservoir algorithm may be recast in the form of surface generation by considering 'accepted' particles as having arrived at the

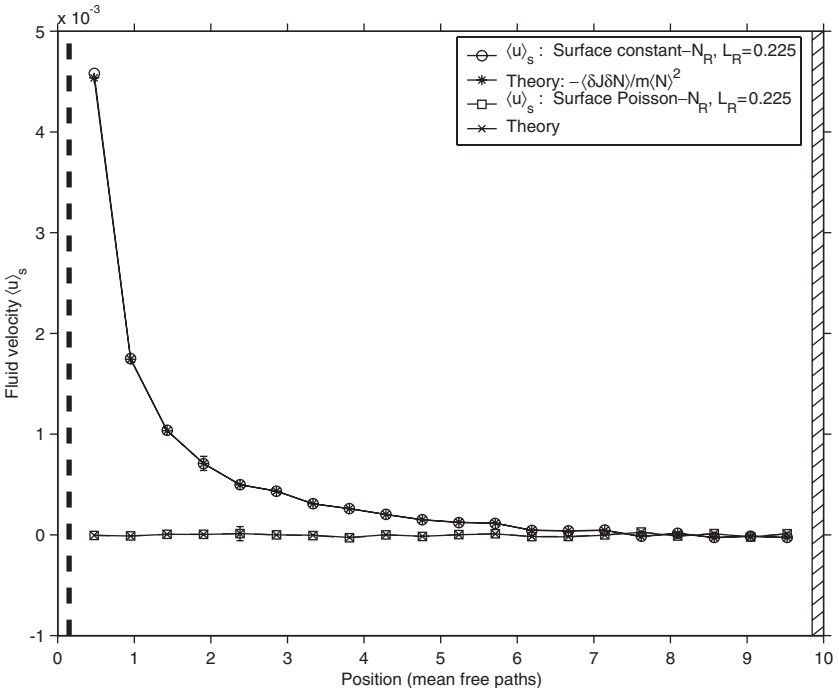

Figure 6. Sample averaging measurement of mean fluid velocity, $\langle u_k^x \rangle_s$, for two variants of the surface generation reservoir algorithms simulating an equilibrium reservoir at temperature $T_R = 5$. The *constant-*$N_{\rightarrow}$ variant (circles) is at non-equilibrium, evidenced by the close correspondence between $\langle u_k^x \rangle_s$ and the $\langle \delta \mathbf{J}\, \delta N \rangle$ correlation (asterisks). The *Poisson-*$N_{\rightarrow}$ variant (squares and x-marks) produces equilibrium behaviour. Statistics were obtained from $S = 5 \times 10^6$ samples.

surface with probability $v^x / v^x_{\max}$. In this case, if the number of reservoir particles to be tested is constant or the number of accepted particles fluxing from the surface is constant then the reservoirs will exhibit a SAM velocity anomaly similar to that described above. However, if either N_R in the first case or $N_{\rightarrow}$ in the second are chosen from a Poisson distribution then the methods should correctly model an equilibrium reservoir.

5. CONCLUDING REMARKS

In this paper we have examined two general reservoir algorithms that are most likely to find practical use in molecular simulations: particle generation in a fixed volume and generation on a surface. In our examples the construction and parameters of both reservoir and system indicate equilibrium conditions, yet we demonstrated that variants of each algorithm induce non-equilibrium behaviour due to non-physical, non-equilibrium correlations of fluctuations. Variants of these algorithms which supply particles to the system at a Poisson-distributed rate, on the other hand, showed no such non-equilibrium behaviour.

Although we have focused on non-physical behaviour in equilibrium reservoir algorithms, we want to emphasize that this is but one example of a more general problem that researchers should be aware of: non-equilibrium behaviour can be computationally introduced into a molecular simulation whenever the inherent stochasticity of fluctuations is not correctly modelled by an algorithm. These effects are subtle (e.g. sample averaging measurement (SAM) velocity in our examples was typically three orders of magnitude smaller than the sound speed) and admittedly unimportant in many cases. However, the effects may become noticeable if physical conditions are extreme or the evolution of the system is sensitive to minor perturbations. For example, if a particle simulation is intended to investigate the onset of an instability [16], an undetected deficiency in a stochastic algorithm could potentially give misleading results. Another example is in simulations of Brownian mechanics [17], where fluctuations are the driving force in the phenomena of interest and so a correct modelling of the stochastic processes is essential. Even when fluctuations are not the topic of interest, researchers working with particle simulations may be puzzled to discover the anomalous SAM velocity when testing their programs by simulating an equilibrium system (which is how the authors discovered this effect).

We have relied extensively on the SAM definition of mean fluid velocity to readily observe anomalous non-equilibrium effects. When an equilibrium state is expected from the system construction, the SAM and CAM definitions should be equivalent; then any statistically significant difference between the SAM and CAM profiles indicates undesirable, computationally induced non-equilibrium conditions. When used this way, the SAM definition of mean fluid velocity is a sensitive tool for detecting non-physical behaviour in a hydrodynamic system. Furthermore, the SAM mean velocity $\langle \mathbf{u} \rangle_s$ is an easy quantity to measure, and a visual comparison of the $\langle \mathbf{u} \rangle_s$ and $\langle \mathbf{u} \rangle_c$ profiles can quickly reveal troublesome areas and assist in problem diagnosis. We expect this tool will find many applications beyond evaluating reservoir performance.

In the Introduction we noted that reservoirs in molecular dynamics (MD) simulations are more complex and require greater finesse in algorithm construction than their counterparts in direct simulation Monte Carlo (DSMC) or lattice gases (LG). In particular, the reservoir can no longer be entirely decoupled from the system; the physically exact nature of movement and collisions in MD mean that certain marginal conditions that are ignored in DSMC or LG must be handled in MD. For example, a particle may have only partly left the reservoir at the end of a time step, so even if it is less than halfway into the system it can still collide with particles in the system. In DSMC and LG, on the other hand, a particle is either inside or outside the system, so this complication does not exist. Thus, in MD simulations the susceptibility of some reservoir algorithms to non-physical non-equilibrium behaviour represents an additional factor to consider in reservoir construction.

REFERENCES

1. Koplik J, Banavar JR. Continuum deductions from molecular hydrodynamics. *Annual Review of Fluid Mechanics* 1995; **27**:257–292.
2. Zheng Y, Garcia AL, Alder BJ. Comparison of kinetic theory and hydrodynamics for Poiseuille flow. *Journal of Statistical Physics* 2002; **109**:495.
3. Macrossan MN. *v*-DSMC: a fast simulation method for rarefied flow. *Journal of Computational Physics* 2001; **173**(2):600–619.
4. Garcia AL, Bell JB, Crutchfield WY, Alder BJ. Adaptive mesh and algorithm refinement using direct simulation Monte Carlo. *Journal of Computational Physics* 1999; **154**:134.

5. Nie X, Chen S, Robbins MO. Hybrid continuum-atomistic simulation of singular corner flow. *Physics of Fluids* 2004; **16**(10):3579–3591.
6. Lilley CR, Macrossan MN. Methods for implementing the stream boundary condition in DSMC computations. *International Journal for Numerical Methods in Fluids* 2003; **42**(12):1363–1371.
7. Hadjiconstantinou N, Garcia AL, Bazant M, He G. Statistical error in particle simulations of hydrodynamic phenomena. *Journal of Computational Physics* 2003; **187**:274–297.
8. Bird GA. *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*. Clarendon Press: Oxford, 1994.
9. Oran ES, Oh CK, Cybyk BZ. Direct simulation Monte Carlo: recent advances and applications. *Annual Review of Fluid Mechanics* 1998; **30**:403–441.
10. Frenkel D, Smit B. *Understanding Molecular Simulation*. Academic Press: New York, 2002.
11. Rivet JP, Boon JP. *Lattice Gas Hydrodynamics*. Cambridge University Press: Cambridge, 2001.
12. Hadjiconstantinou NG. Combining atomistic and continuum simulations of contact-line motion. *Physical Review E* 1999; **59**(2):2475–2478.
13. Delgado-Buscalioni R, Coveney PV. Continuum-particle hybrid coupling for mass, momentum, and energy transfers in unsteady fluid flow. *Physical Review E* 2003; **67**:046704.
14. Tysanner M, Garcia AL. Measurement bias of fluid velocity in molecular simulations. *Journal of Computational Physics* 2004; **196**:173–183.
15. Landau L, Lifshitz E. *Statistical Physics* (Part 1). Butterworth-Heinemann: Stoneham, MA, 1980.
16. Moro E. Hybrid method for simulating front propagation in reaction–diffusion systems. *Physical Review E* 2004; **69**:060101.
17. Reimann P. Brownian motors: noisy transport far from equilibrium. *Physics Reports* 2002; **361**(2–4):257–265.

