INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS
Int. J. Numer. Meth. Fluids 2010; **62**:1081–1106
Published online 8 April 2009 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/fld.2053

Coupled Navier–Stokes—Molecular dynamics simulations using a multi-physics flow simulation framework

R. Steijl and George N. Barakos*, †

Computational Fluid Dynamics Laboratory, Department of Engineering, University of Liverpool, Liverpool L69 3GH, U.K.

SUMMARY

Simulation of nano-scale channel flows using a coupled Navier–Stokes/Molecular Dynamics (MD) method is presented. The flow cases serve as examples of the application of a multi-physics computational framework put forward in this work. The framework employs a set of (partially) overlapping sub-domains in which different levels of physical modelling are used to describe the flow. This way, numerical simulations based on the Navier–Stokes equations can be extended to flows in which the continuum and/or Newtonian flow assumptions break down in regions of the domain, by locally increasing the level of detail in the model. Then, the use of multiple levels of physical modelling can reduce the overall computational cost for a given level of fidelity. The present work describes the structure of a parallel computational framework for such simulations, including details of a Navier–Stokes/MD coupling, the convergence behaviour of coupled simulations as well as the parallel implementation. For the cases considered here, micro-scale MD problems are constructed to provide viscous stresses for the Navier–Stokes equations. The first problem is the planar Poiseuille flow, for which the viscous fluxes on each cell face in the finite-volume discretization are evaluated using MD. The second example deals with fully developed three-dimensional channel flow, with molecular level modelling of the shear stresses in a group of cells in the domain corners. An important aspect in using shear stresses evaluated with MD in Navier–Stokes simulations is the scatter in the data due to the sampling of a finite ensemble over a limited interval. In the coupled simulations, this prevents the convergence of the system in terms of the reduction of the norm of the residual vector of the finite-volume discretization of the macro-domain. Solutions to this problem are discussed in the present work, along with an analysis of the effect of number of realizations and sample duration. The averaging of the apparent viscosity for each cell face, i.e. the ratio of the shear stress predicted from MD and the imposed velocity gradient, over a number of macro-scale time steps is shown to be a simple but effective method to reach a good level of convergence of the coupled system. Finally, the parallel efficiency of the developed method is demonstrated. Copyright © 2009 John Wiley & Sons, Ltd.

Received 16 July 2008; Revised 17 February 2009; Accepted 23 February 2009

KEY WORDS: CFD; mesh-based methods; particle method; object-orientation; multi-physics; Molecular Dynamics; convergence analysis; computational framework; parallel computing

*Correspondence to: George N. Barakos, Department of Engineering, University of Liverpool, Harrison Hughes Building, Liverpool L69 3GH, U.K.
†E-mail: g.barakos@liverpool.ac.uk

1. INTRODUCTION

In recent years, computational fluid dynamics methods based on the (Reynolds-averaged) Navier–Stokes equations have reached a high level of maturity and have been successfully applied to a vast number of flow problems, ranging from academic test cases involving simple geometries to flows around realistic aircraft configurations. The successful application of Navier–Stokes-based flow simulations naturally depends on the validity of the underlying assumptions, i.e. that the flow is continuum and Newtonian. Flow problems for which the continuum assumption breaks down either involve a very low density of the fluid or nano-scale length scales. The validity of the Newtonian assumption depends on the fluid considered, while the length scale of the flow plays a secondary role. For most gas flows, the Newtonian assumption is valid for continuum flow conditions. The breakdown of continuum flow forces one to use either a set of governing equations at a lower level (i.e. molecular level) or extensions to the Navier–Stokes equations valid in the transitional regime, e.g. Burnett equations, along with appropriate boundary conditions. Non-Newtonian behaviour can often be modelled using a modification to the shear stress formulation in the Navier–Stokes equations.

For complex flows in which the continuum and/or the Newtonian assumption breaks down, it is often the case that this breakdown is a gradual process occurring in a limited extend of the considered domain. In aerospace applications, continuum breakdown will occur in low-density flows, e.g. the flow around a spacecraft in the upper layers of the atmosphere. In these conditions, the non-continuum flow features are often limited to the onset of a finite slip velocity on solid surfaces as well as a moderate number of regions of the flow with small-scale flow features, such as the flow through a flap gap or around the wing tips. The majority of the flow field could still be modelled as a continuum flow.

The present work aims to develop a computational framework for flow simulations in which different levels of modelling are used in different parts of the considered domain. The motivation is to extend the range of applications that can be handled, relative to a method based on the Navier–Stokes equations. In addition, the aim is to reduce the computational cost for a given level of fidelity of the modelling by using a less detailed, less expensive model, in parts of the computational domain where this is appropriate. The required coupling of solutions governed by different sets of governing equations constitutes a major challenge in the development of such a computational method. The present paper describes the first steps of the development and application of a framework for such 'multi-physics' flow simulations. The application considered here is the coupling of flow governed by the Navier–Stokes equations with Molecular Dynamics (MD) modelling.

For nano-scale flows in generic geometries, the MD method, described in detail in Refs. [1, 2], can be used directly as a flow simulation method. Examples of such applications can be found in Refs. [3–7]. The coupling of the Navier–Stokes equations for incompressible flow with MD was considered in detail by Ren and Weinan [8]. The work reports on the development and application of their heterogeneous multi-scale method, with the aim to study macro-scale dynamics of fluids in situations, where either the constitutive relation or the boundary conditions are not explicitly available and have to be inferred from microscopic models. The described numerical method builds on the concept of creating multiple micro-scale domains overlapping a macro-scale domain. In their work, continuum hydrodynamics is used as the macro-scale model, while MD serves as the microscopic model and is used to supply the necessary data, for example, shear stress or boundary conditions, for the macroscopic model. Applications were presented for complex fluids, contact-line

DOI: 10.1002/fld

dynamics and fluid–solid interaction. Closely related to the present work is an example of pressure-driven two-dimensional (homogeneous in x-direction) channel flow, in which the momentum equation in x-direction is discretized using central differences and integrated explicitly in time. In the momentum equation, the shear stresses are computed from two-dimensional MD simulations for a Lennard-Jones fluid [9] with an imposed linear velocity gradient. Similar results are also presented for problems, in which the Lennard-Jones particles where replaced with dumb-bells, as a simple model for polymer flow through a channel. Another application of the method was the driven cavity problem, in which the singularity of the stress in the corners, as found in continuum Newtonian models, is handled by using a local MD simulation. The shear stress singularity in the driven cavity problem has been addressed with MD simulation by various researchers. In Ref. [10], MD simulations of the driven cavity were presented and used to investigate the velocity slip and Newtonian behaviour of the shear stresses in the vicinity of the moving wall/side wall corners. An expression for the velocity slip was derived as function of the shear rate, which can be used in a continuum simulation of the flows. The shear stresses were reported to be Newtonian, in contrast to earlier findings of Ref. [4], where a non-Newtonian behaviour was reported for the conditions considered. Nie *et al.* [11, 12] developed a coupled Navier–Stokes/MD method in which a continuum domain partially overlaps with the MD domain(s). In the overlap region(s), the two different flow descriptions provide velocity boundary conditions for the other. A mean atomic flux equal to the continuum flux is imposed at the outer boundary of the atomistic region. In Ref. [11], the coupled method is applied to Couette flow and the flow over a flat plate with nano-scale roughness, while the singular corner problem in driven cavity flows is considered in Ref. [12]. Since this method deals with coupling convective fluxes in the continuum and molecular domains, the method can be regarded as a more direct coupling than the method described in Ref. [8].

In the present work, the Navier–Stokes equations for incompressible flow are coupled to the MD method through the viscous stresses in the discretization of the continuum macro-scale flow model. For the considered test cases, i.e. the Poiseuille flow and the three-dimensional fully channel flow, the flow is homogeneous in x-direction. Therefore, the velocity derivatives in the x-direction and the two cross-flow velocity components v and w vanish. The Navier–Stokes equations for incompressible flow for the considered test cases reduce to

$$\frac{\mathrm{d}u}{\mathrm{d}t} = -\frac{1}{\rho}\frac{\mathrm{d}p}{\mathrm{d}x} + \nu\left(\frac{\mathrm{d}^2u}{\mathrm{d}y^2} + \frac{\mathrm{d}^2u}{\mathrm{d}z^2}\right), \quad \tau_{xy} = \mu\frac{\mathrm{d}u}{\mathrm{d}y}, \quad \tau_{xz} = \mu\frac{\mathrm{d}u}{\mathrm{d}z}$$

Assuming the flow is steady, and relaxing the Newtonian flow assumption, this can be written as

$$\frac{\mathrm{d}p}{\mathrm{d}x} = \frac{\mathrm{d}}{\mathrm{d}y}\tau_{xy} + \frac{\mathrm{d}}{\mathrm{d}z}\tau_{xz} \tag{1}$$

where τ_{xy} and τ_{xz} are the shear stresses in the cross-flow plane. For the Poiseuille flow test case, the y-derivate is omitted. In the proposed hybrid method, the Newtonian shear stresses at selected faces in the finite-volume discretization are replaced with shear stresses predicted by MD simulations for micro-scale domains with the cell face velocity gradient imposed through the Lees–Edwards boundary conditions. Equation (1) is solved using under-relaxed Newton relaxation starting from zero-velocity flow. The Jacobian matrix used in the relaxation is the one derived from Equation (1) assuming Newtonian flow.

DOI: 10.1002/fld

The test cases considered in the present work are similar to those in Ref. [8] in the sense that the coupling decouples the time-scales in the macro-scale continuum domain and the micro-scale MD domains. The present work extends the previous work by addressing in detail the role of sampling errors present in the data from the molecular modelling in the convergence of the coupled system to a steady, repeatable solution. The dependence on the sample duration as well as the number of independent realizations are discussed, along with the methods to reduce the adverse effects of the sampling errors on the convergence.

The design and implementation details of the parallel framework for multi-physics simulations are described in Section 2. The work presented here follows previous descriptions of the multi-physics framework [13, 14]. Based on this framework, the finite-volume method for the incompressible Navier–Stokes equations and the MD method forming the basis for the coupled method are described in Section 3. The application of the coupled method to the planar Poiseuille flow is discussed in Section 4, followed by a further investigation of the convergence behaviour in Section 5. The coupled simulation of the fully developed square channel flow is presented in Section 6. A discussion of the cost effectiveness of the proposed method as well as the parallel efficiency are presented in Section 7. Finally, conclusions are drawn in Section 8.

2. DESIGN OF PARALLEL COMPUTATIONAL FRAMEWORK

The aim of the present work is to develop a framework capable of rapid prototyping of different numerical methods, to facilitate the investigation of coupling strategies between them. The design and implementation of the present computational framework build on a general formulation of the domain decomposition approach. The spatial decomposition idea is widely used in the parallelization of mesh-based as well as particle-based numerical methods. The aim here is to use the commonalities in the implementation of different methods based on spatial decomposition to derive a common set of base classes. The implementation is built as a C++ template class library. The implementation of a numerical method then involves the definition of additional classes, often derived from the library classes, which define the particular details of the method. Similarly, user-defined classes are used to define the spatial decomposition in multiple sub-domains as well as the coupling of the solution in the sub-domains using the concept of 'halo' cells. Message-passing routines are provided by the template class library, building on Message Passing Interface (MPI) functions. The use of halo cells is one of the common features of many numerical methods using spatial decomposition. The coupling of the solution in the various sub-domains takes place by setting the halo cell data based on the solution in the neighbouring sub-domain. This approach is standard in most finite-volume/element methods. In these methods, halo cell values are set either through a direct copy of data internal to the neighbouring domain or through interpolation of the data within the neighbouring domain. The parallel implementation of particle-based methods is typically more demanding due to the dynamic nature of the particle distribution during a simulation. Details on the parallelization of different MD methods can be found in Refs. [13, 15–20], while examples of the use of object-orientation techniques in the implementation of flow simulation methods can be found in Refs. [19, 21–23]. Three parallelization methods have found wide-spread application: spatial decomposition, particle decomposition and 'force' decomposition. The first approach is used in the present work and is based on the formation of a background mesh, in

DOI: 10.1002/fld

which each of the cells holds a variable number of particles. Similar to the mesh-based methods, the background mesh with the associated particles is decomposed into a number of sub-domains, which are distributed over the processors. The commonality with mesh-based method consists of the fact that halo cells are defined around each of the sub-domains holding copies of particles from the neighbouring sub-domain. In this case, a halo cell update constitutes a copy of particles from the neighbouring domain. For cases without a significant disparity in particle density, spatial decomposition provides an effective method for parallel execution. The particle and force decomposition amount to distributing the particle and particle–particle interactions across the processors, respectively. Depending on the physical problem considered, these approaches may be more efficient than the spatial decomposition approach, for example when a relatively small number of macro-molecules are considered in a MD simulation. In addition to the commonality in halo cell data exchange, the present method provides implementation for a number of algorithms commonly used in numerical methods. Particle-based simulation methods for different sets of governing equations often share common algorithms, examples of which are discussed in Ref. [24]. An example used in the present work is the cell linked-list algorithm, illustrated in Figure 1, for MD with short-range interactions defined through the Lennard-Jones inter-particle potential. The details of these methods are given in the following paragraphs. A background mesh is defined providing cells with a size exceeding the interaction range of the particles. Depending on the particle coordinates, each particle is binned in one of the cells. The inter-particle potential computation then requires a search through the cell in which the particle is located as well as its 26 neighbour cells.

Figure 2 shows an example of spatial decomposition of an MD domain. The halo cells for the two sub-domains as well as the data communication are shown in Figure 2(a). The present framework provides an algorithm to construct for each domain face, a send/receive data buffer into which the sending process copies the particles located in the layer of cells along the domain boundary. The receiving process then constructs copies of these particles based on the received data. This process is sketched in Figure 2(b). In the course of a MD simulation, particles will move from one cell to another and, hence less frequently, from one sub-domain to a neighbouring domain. The algorithm implemented for this particle transfer is sketched in Figure 2(c). Similar to

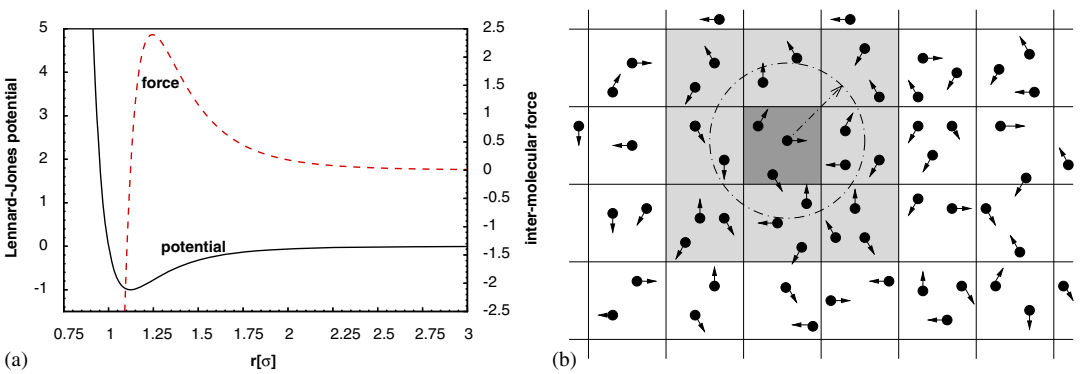

Figure 1. (a) Lennard-Jones inter-molecular potential and (b) cell linked-list algorithm employed in MD simulations with short-range interaction potential.

DOI: 10.1002/fld

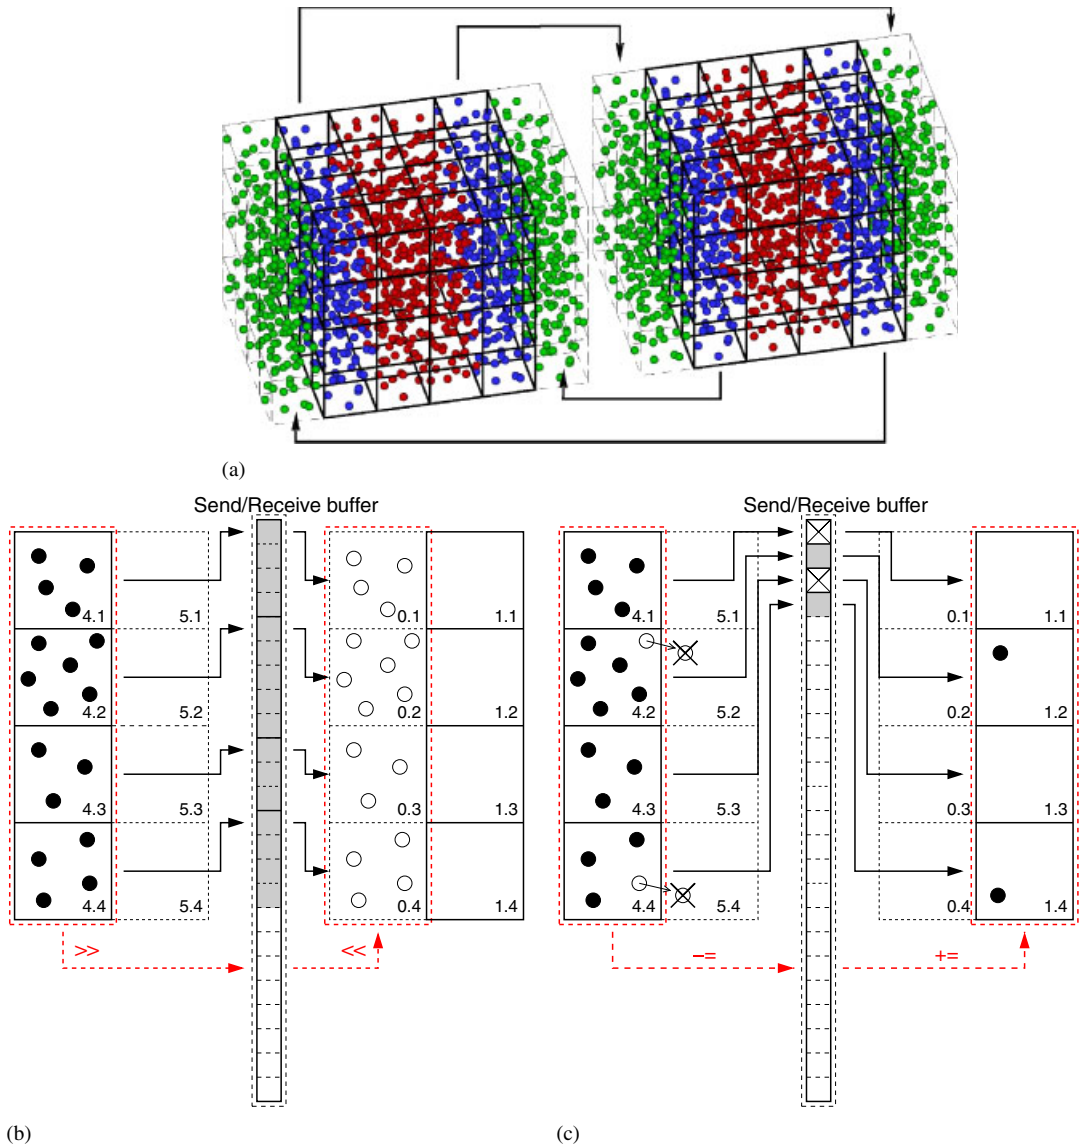

Figure 2. Parallel Molecular Dynamics method based on spatial decomposition: (a) associations of particles with halo cells of neighbouring sub-domains and (b) translation of particle from sub-domain to a neighbouring sub-domain. For both operations, the implementation uses operator overloading as well as the formation of send/receive buffers and message-passing to communicate the buffer data.

the halo particle exchange, a data buffer (however, for this exchange, much smaller) is constructed, which will hold the particles leaving the domain on the sending process. The receiving process then creates a new particle in the interior of the neighbouring domain based on the received data.

DOI: 10.1002/fld

3. HYBRID FINITE-VOLUME/MD METHOD

To demonstrate the current framework, a standard MD with short-range interaction potential was implemented and coupled to a finite-volume method for the Navier–Stokes equations for incompressible flow.

3.1. MD method

The MD method [1] simulates the dynamics of a system of N interacting molecules by temporal integration of Newton's equations of motion:

$$m_i \frac{\mathrm{d}^2 \underline{x}_i}{\mathrm{d}t^2} = F_i, \quad i = 1, \ldots, N \tag{2}$$

where $\underline{x}_i$ denote the particle positions and the forces F_i are the negative derivatives of a potential function $U_i(\underline{x}_1, \ldots, \underline{x}_N)$. In the present work, the system of N particles is integrated in time from level n to $n+1$ using the velocity form of the second-order accurate, time reversible, Verlet algorithm [25]:

$$\begin{aligned} \underline{x}_i^{(n+1)} &= \underline{x}_i^{(n)} + \delta t \underline{v}_i^{(n)} - \frac{\delta t^2}{2m_i} \nabla U_i^{(n)}, \quad i = 1, \ldots, N \\ \underline{v}_i^{(n+1)} &= \underline{v}_i^{(n)} + \frac{\delta t}{2m_i} \nabla U_i^{(n)}, \quad i = 1, \ldots, N \end{aligned} \tag{3}$$

where $\underline{v}_i$ denotes the velocity of particle i. The Verlet algorithm is a standard temporal integration method in MD. The use of Newton's equation of motion automatically implies the use of classical mechanics to describe the motion of the atoms. At normal temperatures and for most atoms, this use of classical mechanics is justified. However, for intra-molecular atomic motions in poly-atomic molecules, e.g. vibrational motions, quantum mechanical effects often need to be included. In the present work, the Lennard-Jones fluid is used as a model, i.e. an atomic medium with an inter-particle potential given by

$$U_{\mathrm{LJ}}(r_{ij}) = 4\varepsilon_{ij} \left[\left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}} \right)^{6} \right] \tag{4}$$

where r_{ij} is the distance between particles i and j, ε_{ij} is the depth of the potential well and σ_{ij} is the (finite) distance at which the inter-particle potential is zero. The Lennard-Jones potential combines strong repulsion at short distances with weak attraction at longer distances. In the above equation, σ_i as well as ε_i are assumed to be identical for $i = 1, \ldots, N$. Furthermore, a cut-off distance r_c is defined. Therefore, particles within this distance interact through Equation (4), while beyond this range, the contribution to the inter-particle potential is neglected. Here, the cut-off distance is chosen such that $r_c = 4.0\sigma$. Long-range interactions, such as Coulomb forces, are not included and the present method only considers short-range particle interactions based on Equation (4). The inter-particle potential and force are shown in Figure 1(a). The implementation of the velocity-Verlet scheme utilizes the cell linked-list algorithm, which divides the computational domain in uniform cubic cells of dimension d, where $d = r_c$ or slightly larger, as illustrated in Figure 1(b) for the two-dimensional case. In a three-dimensional simulation, this reduces the task of finding

DOI: 10.1002/fld

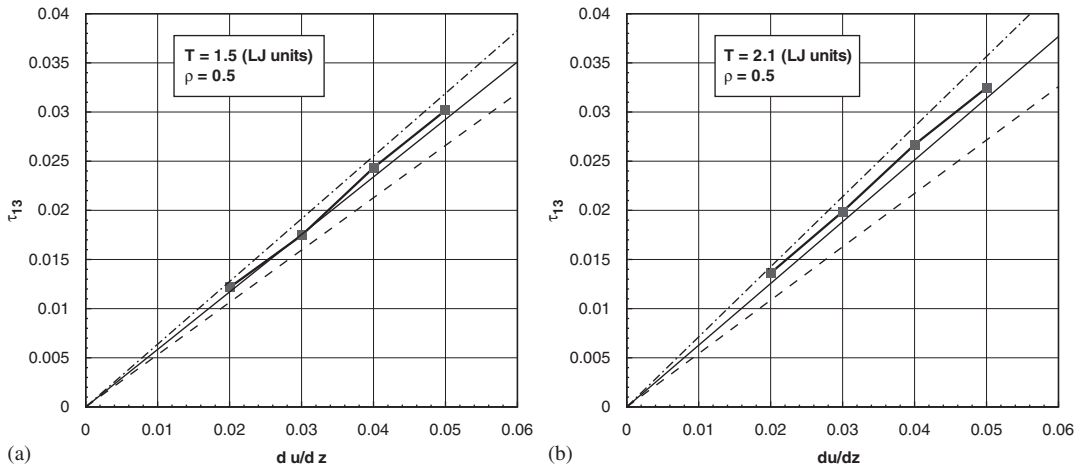

Figure 3. Prediction of shear stress as function of shear rate from Molecular Dynamics simulation with Lees–Edwards boundary condition. The predicted results for two temperatures are compared with the representative data from the literature [27]. The dashed and the dash-dot lines denote the upper and lower bounds of the compiled literature data: (a) $\rho = 0.5\sigma^{-3}$, $T = 1.5\varepsilon/k_B$ and (b) $\rho = 0.5\sigma^{-3}$, $T = 1.5\varepsilon/k_B$.

neighbours of a given particle to checks over 27 cells, i.e. the cell in which the cell is located and its 26 nearest neighbour cells.

3.2. Validation of MD method for stress prediction

In the present work, MD simulations are used as a means of evaluating shear stresses as function of imposed velocity gradients. Based on the computational framework described in Section 2, an MD method using spatial decomposition was developed. Lees–Edward boundary conditions [26] are used to impose linear velocity gradients. Figure 3 shows the predicted shear stresses as function of shear rate from the developed method. In Lennard-Jones units, the density of the fluid is $0.50\sigma^{-3}$, while the temperature is held constant at $1.50\varepsilon k_B^{-1}$ and $2.10\varepsilon k_B^{-1}$, for the two considered cases, respectively. Here k_B is the Boltzmann constant. In Lennard-Jones units, time is normalized by $(\varepsilon/\sigma^2)^{-1/2}$, while the shear stress and the shear rate have dimensions $\varepsilon\sigma^{-3}$ and $(\varepsilon/\sigma^2)^{1/2}$, respectively. In the following, an LJ time units is denoted by τ. The Berendsen thermostat [28] is applied to the thermal velocities of the fluid to control the temperature. A cubic domain was used with a 10^3-cell linked-list and each cell is 4.0σ in each direction. This cut-off was used to obtain reliable viscosity evaluation from the Irving–Kirkwood relation [29]. The shown results were obtained after averaging over 200 000 time steps of 0.001 LJ time units, following an initial equilibration phase of 50 000 time steps. The predicted results for two temperatures are compared to representative data from the literature [27]. The dashed and the dash-dot lines denote the upper and lower bounds of the literature data. The difference in the reported values of the shear viscosity for the Lennard-Jones fluid model shows that determining shear stresses and shear viscosities from MD is still a challenging problem, also reflected by the number of recent works in the literature dealing with this topic [27, 30–33].

DOI: 10.1002/fld

3.3. Finite-volume discretization

For the two test cases considered here, i.e. planar Poiseuille flow and fully developed flow in a square duct, the flow is assumed steady incompressible with homogeneous flow in x-direction. A constant linear pressure gradient is assumed in the x-direction. The velocity components in y- and z-directions vanish under these assumptions. The flow simulation therefore reduces to solving for the velocity component u in x-direction from the momentum equation in x-direction. For the Poiseuille problem, the finite-volume mesh is sketched in Figure 4. The problem is solved using a cell-centred method. For cell k, numerical fluxes through cell face k and $k+1$ are required. Two flux computation methods are used; the first one uses a simple central difference evaluation of the velocity gradient normal to the cell face. For a uniform mesh, this scheme is second-order accurate in space. In addition, a third-order accurate flux reconstruction method is considered. This method computes the derivative in cell face k from cell centres $k-2, k-1$, k and $k+1$ and for the cell faces in the vicinity of the boundaries, asymmetric stencils are used. The third-order method is derived for a stretched Cartesian mesh. The two-dimensional mesh for the fully developed channel flow is sketched in Figure 5. The spatial discretization for this case is based on a cell-centred finite-volume method, formulated for a general curvilinear mesh, using central differences for the evaluation of the velocity gradients in the calculation of the fluxes in the y- and z-directions.

4. PLANAR POISEUILLE FLOW

The first test case considered in the current setup was the planar Poiseuille flow. The x-axis is in the homogeneous direction and the z-axis points in the cross-flow direction. In Lennard-Jones units [1, 2], the domain height is 20σ. The Lennard-Jones potential is particularly representative for noble gases, such as Argon, in which case $\sigma = 3.4 \cdot 10^{-10}$ m. Consequently, the considered channel size is at the nano-meter scale. The density and temperature of the fluid were $0.80\sigma^{-3}$ and $1.50\varepsilon k_B^{-1}$, respectively, and a constant pressure gradient of $0.02\varepsilon\sigma^{-4}$ in the x-direction is imposed. The viscous stresses on the cell faces are evaluated using micro-scale MD simulation with the $\mathrm{d}u/\mathrm{d}z$ gradient of the cell face imposed on the problem, as sketched in Figure 4. For the 8-cell uniform mesh considered here, the symmetry of the problem leads to four different cell face gradients, for which a micro-scale problem is constructed to evaluate the shear stress. For each of the MD problems, a cubic domain with dimension 12σ is defined with 1382 particles, located in a 3^3 cell linked-list with a cut-off length of 4.0σ. Temporal integration using the velocity-Verlet method [25] with time step 0.001τ is used. The Berendsen thermostat [28] is applied to the thermal velocities of the fluid to control the temperature within the MD domains. The velocity gradients are imposed using Lees–Edwards boundary conditions [26]. Each of the micro-scale problems was first run for 50 000 time steps to reach equilibrium. Subsequently, the shear stresses evaluated with the Irving–Kirkwood relation [29] were averaged over 200 000 time steps. For each micro-scale problem, the average of either 5 or 10 independent realizations is used.

For the Poiseuille problem considered here, the finite-volume discretization of the momentum equation in x-direction leads to the following equation for the residual vector $\underline{r}$:

$$r_k = V_k \frac{\mathrm{d}p}{\mathrm{d}x} - [A_{k+1}(\tau_{xz})_{k+1} - A_k(\tau_{xz})_k] \tag{5}$$

DOI: 10.1002/fld

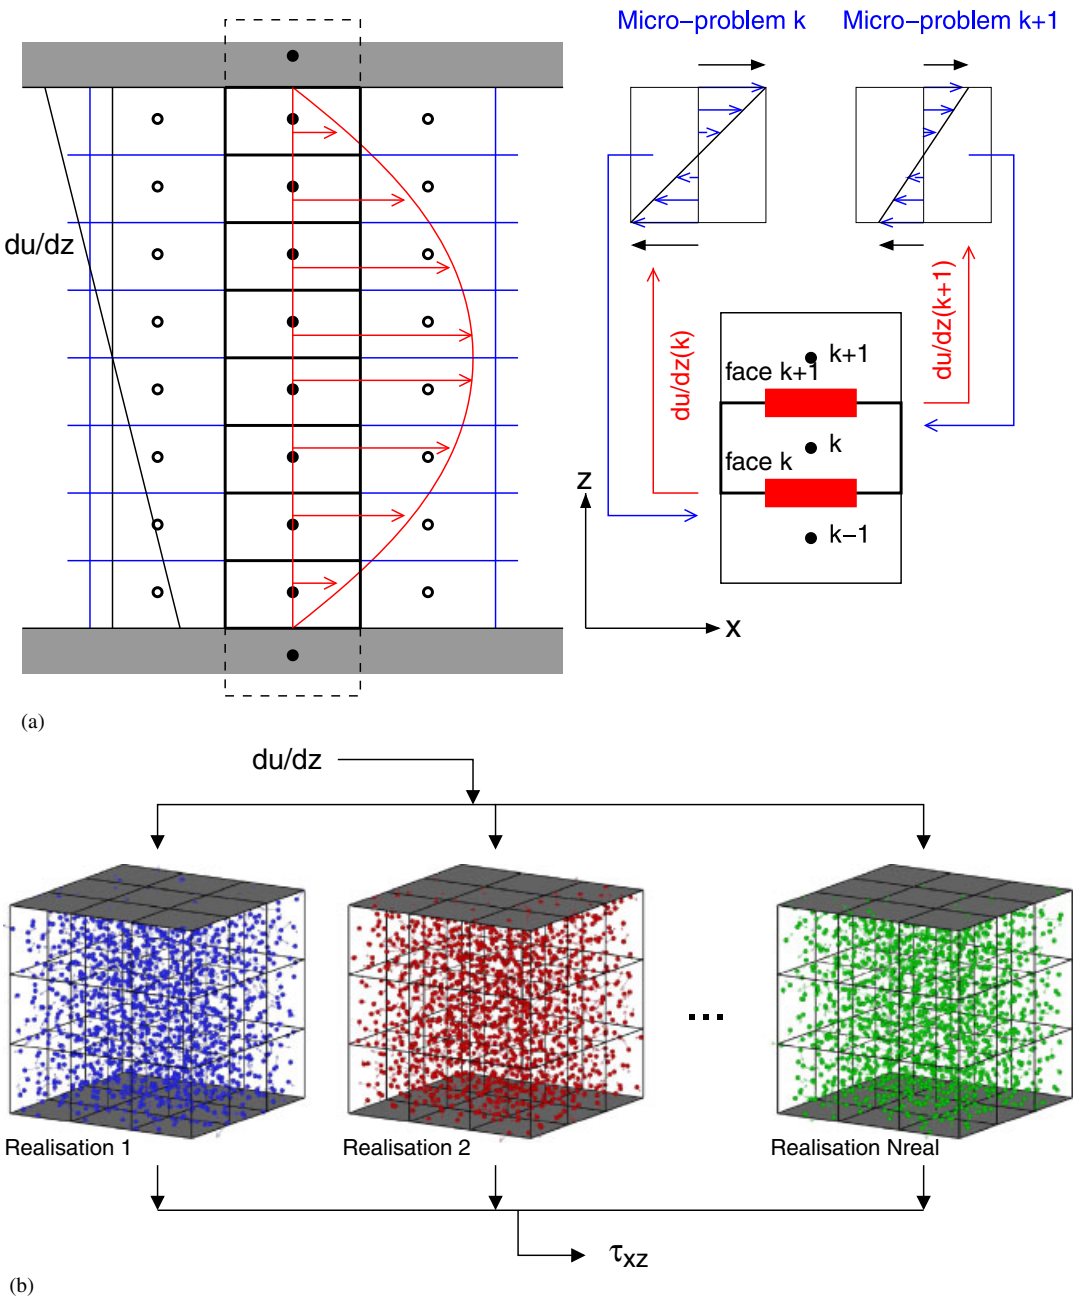

Figure 4. Coupled Navier–Stokes/Molecular Dynamics method for Poiseuille flow: (a) cell-centred finite-volume method, numbering of unknowns and cell faces and (b) Molecular Dynamics method with Lees–Edwards boundary conditions for imposed velocity gradient in one direction.

DOI: 10.1002/fld

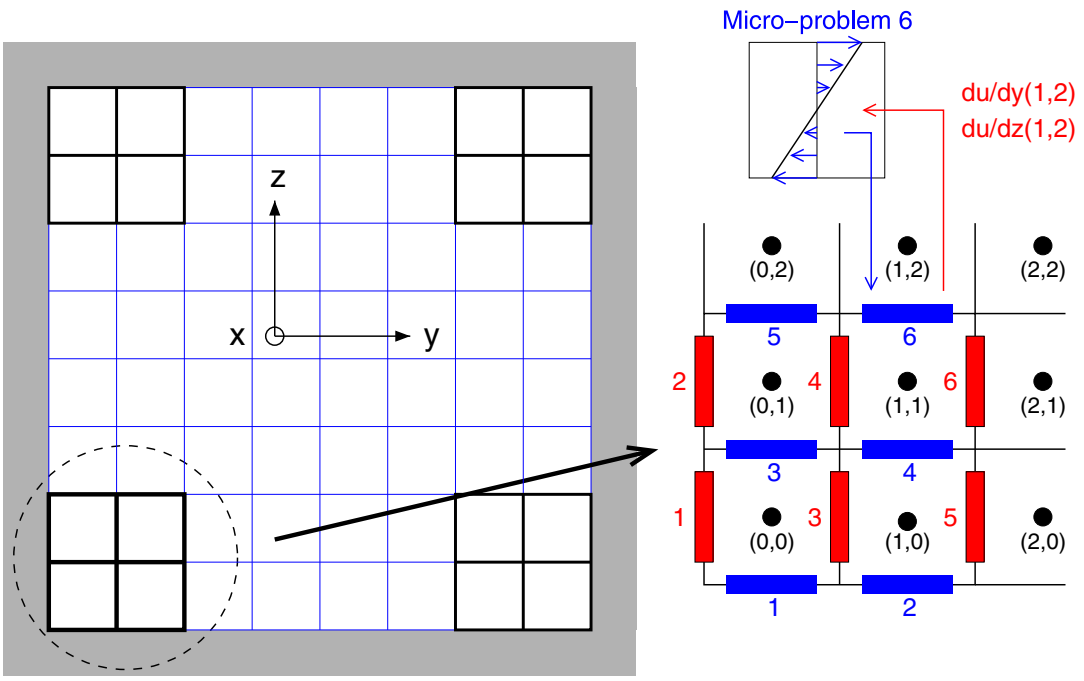

(a)

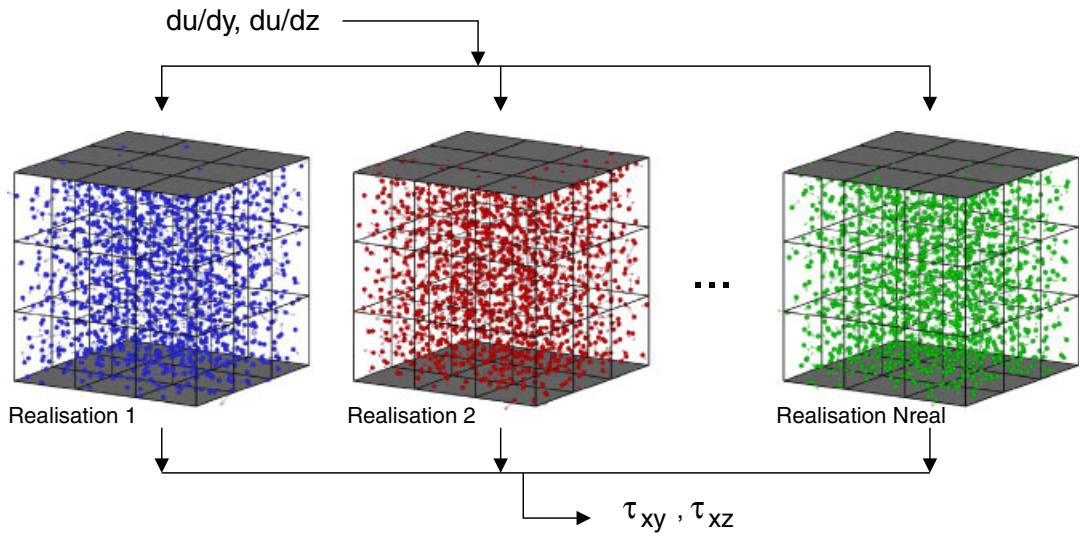

(b)

Figure 5. Coupled Navier–Stokes/Molecular Dynamics method for fully developed 3D channel flow, with flow in x-direction: (a) cell-centred finite-volume method, placement of micro-scale problems in channel corners and numbering of cell faces and (b) Molecular Dynamics method with Lees–Edwards boundary conditions for imposed velocity gradient in two directions.

DOI: 10.1002/fld

where A_k represent the cell face area of cell k in the (x, y) plane and V_k is the volume of cell k. An under-relaxed Newton iteration is used to solve for the velocity vector $\underline{u}$,

$$\underline{u}^{v+1} = \underline{u}^{v} - \omega \mathbf{A}^{-1} \underline{r} \tag{6}$$

where $\mathbf{A}$ represents the flux Jacobian matrix $\mathrm{d}\underline{r}/\mathrm{d}\underline{u}$ and ω is the relaxation factor. To facilitate the construction of the Jacobian matrix, the assumption of Newtonian flow is made in the evaluation of the partial derivatives in the Jacobian matrix elements. Hence, the Jacobian matrix is an approximation of the exact Jacobian of the coupled system. The second-order accurate discretization on a uniform mesh leads to a tri-diagonal matrix, while the third-order accurate stretched mesh discretization leads to a penta-diagonal matrix.

The simulation was initialized using an exact solution based on the assumption of Newtonian flow with constant viscosity, re-scaled with a factor of 0.75 to ensure that the numerical solution of the coupled system involves significant updates of the velocity. The evolution of the computed velocity as a function of the Newton step is shown in Figure 6 for two of four independent simulations computed for the coupled problem. The results indicate that for the present problem, the velocity field approaches the exact solution within about five Newton relaxation steps. Figure 7 shows the convergence of the L_2 norm of the residual vector of the macro-scale finite-volume method. It is apparent that using the MD predictions for the shear stresses on the cell faces directly leads to a stalled convergence after about five Newton steps (Figure 7(a)). At this stage, the major updates to the velocity field had occurred, and subsequent relaxation steps will have velocity gradients at the cell faces that are increasingly closer to the gradient during the preceding relaxation step. The role of the statistical scatter in the computed MD shear stresses then becomes more prominent. The results of Figure 7 indicate that sampling the stresses in 10 independent realizations does not lead to significantly better convergence than sampling of just 5 realizations.

To illustrate this statistical scatter, Figure 8 shows the convergence of the shear–stress statistics in the micro-scale MD problems for a representative Newton relaxation step of the coupled Navier–Stokes/MD simulation. It can be seen that after the initial velocity updates have occurred, the updates in the velocity gradients become of similar order as the scatter in the predicted shear for a given imposed velocity gradient on the MD problem. In the results shown, sampling of the statistics starts after 50 000 MD time steps. Solid lines represent the average of five independent realizations and error bars denote the variability of the statistics within the group of these five realizations.

A simple method to reduce the effects of statistical scatter is to replace the MD shear stress prediction at face k for Newton relaxation step v with a reconstruction based on the averaging of an 'apparent' viscosity over the preceding relaxation steps, i.e.

$$(\tau_{13})_k^v = \left[\frac{1}{v - v_{\mathrm{start}} + 1} \sum_{i=v_{\mathrm{start}}}^{v} \frac{(\tau_{13})_k^i}{\left(\dfrac{\mathrm{d}u}{\mathrm{d}z}\right)_k^i} \right]_k \left(\frac{\mathrm{d}u}{\mathrm{d}z}\right)_k^v \tag{7}$$

where the term between square brackets denotes the averaged 'apparent' viscosity and v_{start} is the iteration step at which the averaging commences. The above process is applied when the velocity field has converged to the stage at which the statistical scatter starts to dominate, in the present example $v_{\mathrm{start}} = 5$. The averaging operation of Equation (7) can be used in the general case of fluids for which the shear viscosity shows a significant dependence on the shear rate, such

DOI: 10.1002/fld

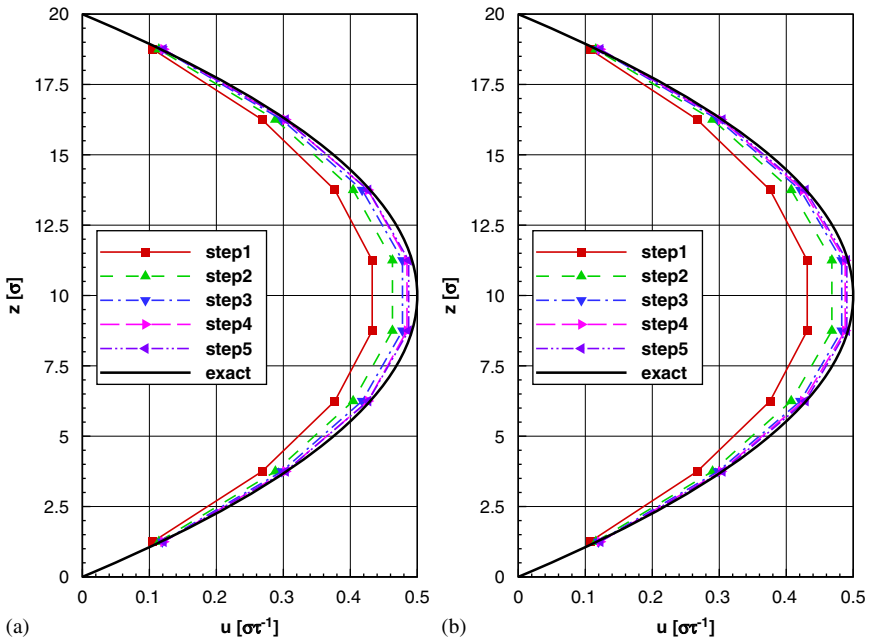

Figure 6. Coupled Navier–Stokes/Molecular Dynamics method for Poiseuille flow. Convergence of velocity profile towards exact solution for estimated viscosity: (a) realisation 1 and (b) realisation 2.

as in shear-thinning fluids, since the initial relaxation phase without averaging will give a good approximation of the local velocity gradient before the averaging commences. The improvement in the convergence of the coupled system is shown in Figure 7(b).

5. CONVERGENCE ANALYSIS FOR POISEUILLE FLOW

The averaging of an 'apparent' viscosity as used in the previous section is analysed in more detail here for a less expensive model problem. The planar Poiseuille problem for Newtonian flow is considered in which the scatter in the MD statistics, as found in the model of the previous section, is represented by a random noise imposed on the otherwise constant viscosity. The residual for a cell k then becomes

$$r_k = V_k \frac{\mathrm{d}p}{\mathrm{d}x} - \mu \left[A_{k+1}(1+\delta\,\mathrm{rand}())_{k+1} \left(\frac{\mathrm{d}u}{\mathrm{d}z}\right)_{k+1} - A_k(1+\delta\,\mathrm{rand}())_k \left(\frac{\mathrm{d}u}{\mathrm{d}z}\right)_k \right] \tag{8}$$

where rand() returns uniformly distributed random numbers with zero mean and δ is the noise amplitude.

Figure 9 shows the convergence as function of Newton relaxation step for the Poiseuille flow with the noise added to the viscosity. Three finite-volume meshes with 40 cells are considered: a uniform mesh, a stretched mesh with a constant expansion ratio of 1.10 and a stretched mesh with a constant expansion ratio of 1.20. Two finite-volume discretization methods are also compared. The first is the second-order central difference method for uniform meshes, as used in the previous

DOI: 10.1002/fld

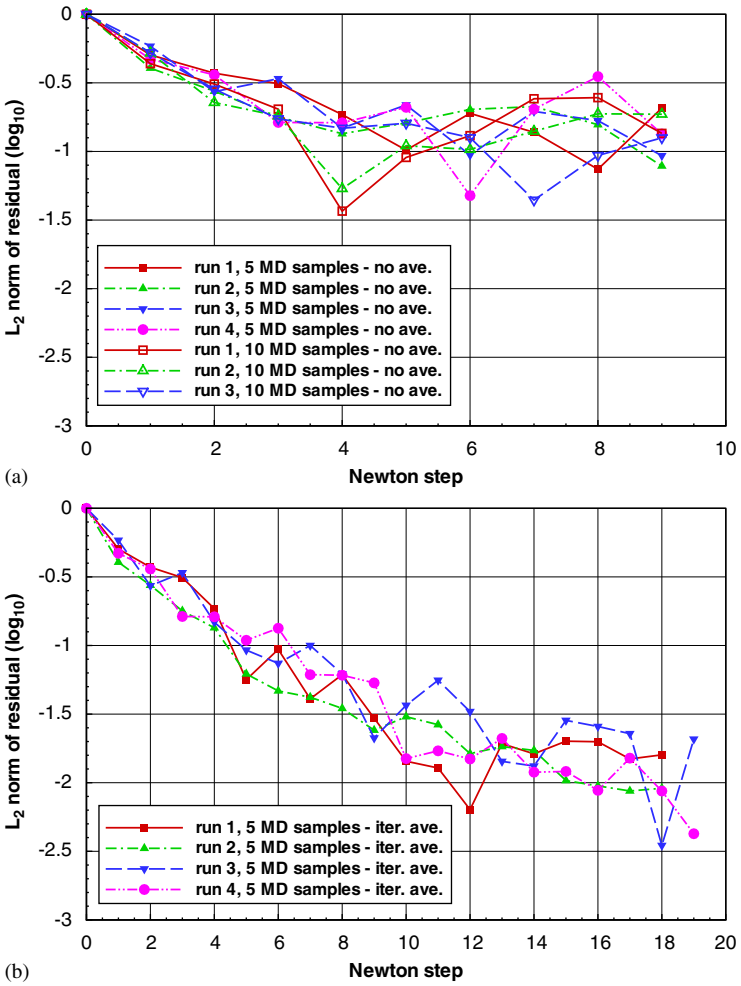

Figure 7. Convergence of the L_2 norm of the residual vector for the coupled Navier–Stokes/Molecular Dynamics simulation of the planar Poiseuille problem: (a) using the MD predictions for the shear stresses on the cell faces directly leads to a stalled convergence and (b) improved convergence with averaging of 'apparent' viscosity over previous iterations, after Newton iteration 5.

section for the coupled Navier–Stokes/MD method. The second method uses a third-order accurate gradient computation on the cell face using two cell-centred values left and right of the cell face. The third-order method is formulated for non-uniform Cartesian meshes. For noise amplitudes δ of 0.2 and 0.4%, the evolution of the maximum residual as well as the L_2 norm of the residual vector are compared. As in the previous section, results without averaging show an apparent lack of convergence. The results shown are taken from 100 independent simulations, to obtain meaningful statistics of the residual convergence. As expected, a larger noise amplitude leads to a smaller reduction of the residual in the macro-domain discretization, both in maximum value as well as is L_2 norm. The averaging of the 'apparent' viscosity leads to a converging solution, for all

DOI: 10.1002/fld

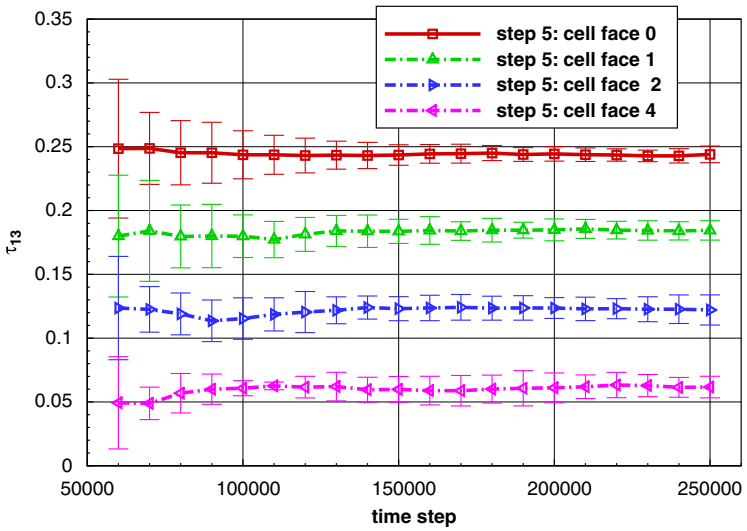

Figure 8. Convergence of shear–stress statistics in Molecular Dynamics micro-scale problem in hybrid Navier–Stokes/Molecular Dynamics method for Poiseuille flow. Sampling of statistics starts after 50 000 MD time steps. Solid lines represent the average of five independent realizations. Error bars denote the variability of the statistics within the group of five realizations.

cases considered, and the results indicate that the third-order accurate discretization is more noise sensitive than the central second-order scheme. Interestingly, the stretching of the mesh does not lead to significant changes in the convergence behaviour.

In Equation (8), the statistical scatter in the shear stresses scales linearly with the shear rate. The uncertainty in shear stresses evaluated with MD simulations does not show this linear dependence; i.e. at lower shear rates, the statistical scatter typically has a larger relative magnitude than at higher shear rate, as is apparent from the published results [27, 31]. The convergence analysis was repeated with a random disturbance amplitude of 0.2 and 0.4% of the wall shear stress added to the shear stress instead of the perturbed viscosity to model the relatively larger scatter at lower shear rates. The results were qualitatively as well as quantitatively similar to those shown in Figure 9. Again, the simulations without averaging show a lack of convergence, while the averaging of the 'apparent' viscosity solves this problem.

6. FULLY DEVELOPED CHANNEL FLOW

As a second test case, a three-dimensional fully developed flow is considered in a square channel. The flow is assumed homogeneous in the x-direction; hence, the problem is reduced to two-dimensional with non-zero-velocity gradients in the y- and z-direction and constant velocity in the x-direction. The pressure is assumed constant in the (y, z) plane. In Lennard-Jones units, the channel is 20σ wide in both directions. The density of the fluid is $0.80\sigma^{-3}$, while the temperature is held constant at $1.50\varepsilon k_{\mathrm{B}}^{-1}$. A constant pressure gradient of $0.05\varepsilon\sigma^{-4}$ in the x-direction is imposed. The macro-scale problem comprises the finite-volume discretization on a stretched Cartesian mesh

DOI: 10.1002/fld

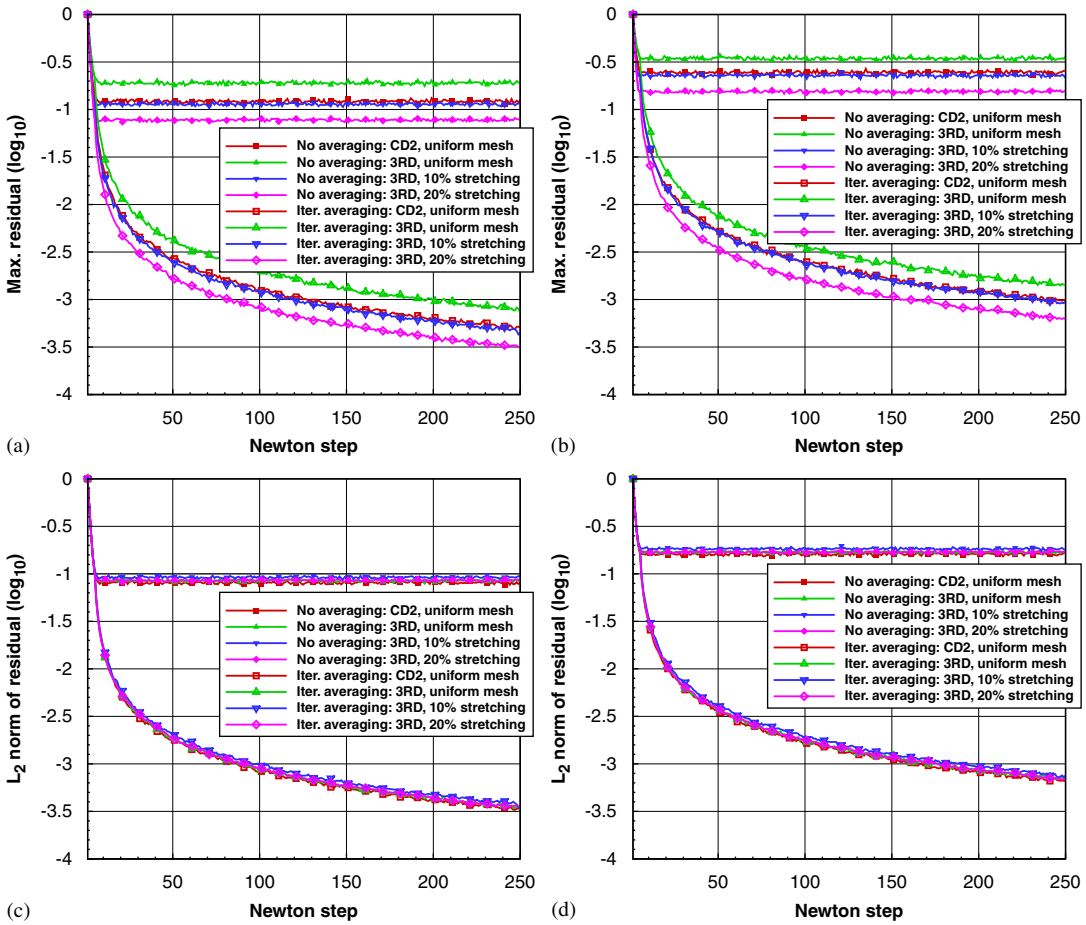

Figure 9. Convergence of Poiseuille flow with noise added to viscosity. Finite-volume discretization with 40 cells, under-relaxed Newton relaxation ($\omega = 0.5$). Results for 0.2 and 0.4% random disturbance on viscosity are compared for different spatial discretizations: (a) maximum residual, 0.2% noise; (b) maximum residual, 0.4% noise (c) L_2 norm, 0.2% noise; and (d) L_2 norm, 0.4% noise.

with 40×40 cells in cross-flow direction of the incompressible Navier–Stokes equations. The mesh is clustered towards the domain boundaries and expands with a constant ratio of 1.10 towards the domain centre. The finite-volume discretization of the momentum equation in x-direction reduces to the following equation for the residual in cell (j,k):

$$r_{j,k} = V_{j,k}\frac{\mathrm{d}p}{\mathrm{d}x} - [(A_{xz})_{j+1,k}(\tau_{xy})_{j+1,k} - (A_{xz})_{j,k}(\tau_{xy})_{j,k}] - [(A_{xy})_{j,k+1}(\tau_{xz})_{j,k+1} - (A_{xy})_{j,k}(\tau_{xz})_{j,k}] \quad (9)$$

where $(A_{xy})_{j,k}$ and $(A_{xz})_{j,k}$ represent the cell face area of cell j,k in the (x,y) and (x,z) planes, respectively, and $V_{j,k}$ is the volume of cell (j,k). A cell-centred finite-volume discretization for

DOI: 10.1002/fld

this problem is used, formulated for a general curvilinear mesh, using central differences for the evaluation of the velocity gradients. The Jacobian matrix for the hybrid system then is a sparse matrix with a maximum of five non-zero entries per row. At each Newton step, the system is solved using a pre-conditioned conjugate gradient method. For a group of four cells in the domain corners, the viscous stresses on the cell faces are evaluated using micro-scale MD simulation with the $\partial u/\partial y$ and $\partial u/\partial z$ gradients of the cell face imposed on the problem, as sketched in Figure 5. Using the symmetry of the problem, six micro-scale problems for the shear stress evaluation need to be solved. For each of the MD problems, a cubic domain with dimension 12σ is defined with 1382 particles, located in a 3^3 cell linked-list with a cut-off length of 4.0σ. Temporal integration using the velocity-Verlet method [25] with a time step 0.001 is used. The Berendsen thermostat [28] is applied to the thermal velocities of the fluid to control the temperature within the MD domains. The velocity gradients are imposed using Lees–Edwards boundary conditions. Each of the micro-scale problems is first run for 50 000 time steps to reach equilibrium. Subsequently, the shear stresses evaluated with the Irving–Kirkwood relation [29] are averaged during a sampling phase of at least 100 000 time steps. In this section, the effect of different sampling methods for the shear stresses predicted by MD is investigated. The first set of data, termed 'baseline' in the following, employs four independent realizations for each of the micro-scale MD problems, while a sampling interval of 100τ (100 000 time steps) is used. Multiple realizations of the coupled Navier–Stokes/MD approach are presented in the current work to highlight variability in the convergence behaviour. For this baseline data set, the averaging of the apparent viscosity is used. In addition to this data set, results were obtained for 8 and 16 independent realizations for each micro-scale MD problem (with a sample duration of 100τ) as well as for four independent realizations for each micro-scale MD problem with an extended sampling interval of 200τ.

Figures 10, 11 show the convergence of the residuals versus Newton relaxation step. Results for the maximum as well as the L_2 norm of the residual vector are presented. The simulations were initialized with zero velocity and subsequently marched through 25 Newton relaxation steps assuming Newtonian flow; i.e. for steps 0–24, the viscous stresses were all evaluated directly from the Newtonian expression. An under-relaxation factor of 0.25 was used. For the cell faces for which the MD micro-problems are used subsequently, the velocity gradients are shown in Table I. The rapid linear reduction of both residual norms in the logarithmic plots of Figures 10, 11 demonstrates the capability of the employed finite-volume method to generate a well-converged numerical solution for the considered problem, i.e. with the employed under-relaxation factor of 0.25 it would take approximately 100 relaxation steps to reach machine precision, as previous experimentation has revealed.

Following this initialization phase, for relaxation step 25 and higher, the coupled Navier–Stokes/MD approach is used. The replacement of the Newtonian shear stresses with the ones from the MD problems evidently lead to an initial increase of the maximum as well as the L_2 norm of the residual vector. In the subsequent 20 relaxation steps, both residual norms can be seen to decrease to varying levels depending on the employed micro-scale sampling approach.

Figure 10 compares the convergence enhancement due to averaging-of-apparent viscosity with that due to doubling the sampling time. Results are compared for two sets of results. The first set employs four realizations per MD micro-problem and a sampling length of 200 LJ time units, while the averaging of apparent viscosity is not used. The second set of results does apply the averaging of apparent viscosity, while the four realizations per MD micro-problem now have a reduced sampling length of only 100 LJ time units. The results with averaging of apparent viscosity clearly show a better convergence in both residual norms. The present averaging-of-apparent viscosity

DOI: 10.1002/fld

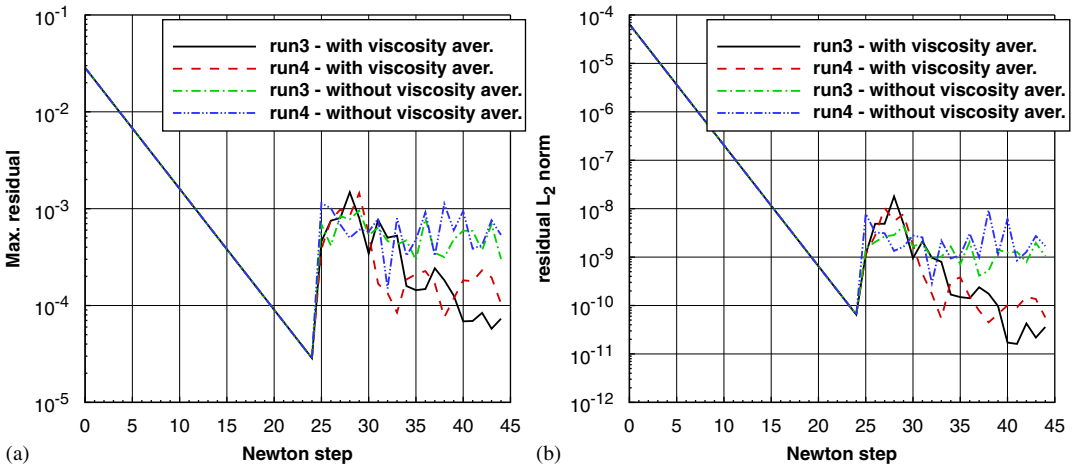

Figure 10. Convergence of 3D channel flow simulation. Finite-volume discretization with 40×40 cells, under-relaxed Newton relaxation ($\omega = 0.25$). The simulations employ four realizations per MD micro-scale problem. The results with averaging-of-apparent viscosity use a 100τ sample duration while a 200τ sampling time is used for the results without averaging: (a) maximum residual and (b) L_2 norm of residual.

method leads to a better convergence of the coupled system than either doubling of the sampling time or increasing the number of independent samples to 8 or even 16. The computational overhead involved with computing the averages of the apparent viscosity in the proposed method comprises a small number of very inexpensive operations relative to simulations without this averaging; i.e. for each cell face, the apparent viscosity needs to be computed from the stress components, followed by a global reduction operation in the case of parallel execution, then the averaging with the average over the previous relaxation steps and, finally, the computation of the viscous stress based on the current velocity gradients. These steps are only required at the end of the MD sampling, so are done only once per Newton relaxation step. The incurred overhead is therefore many orders of magnitude smaller than the cost of a single MD simulation.

Figure 11 shows results for four independent simulations for the coupled problem using the 'baseline' sampling. These results are compared with the results obtained with 8 or 16 realizations per MD micro-problem (while the sampling length was kept at 100 LJ time units). In addition, the set of results with 'improved' sampling includes results for cases with four realizations per MD micro-problem and an extended sampling time of 200 LJ time units. For the first five Newton steps of the coupled problem (i.e. steps 25–29 in the figure), when the averaging-of-apparent-viscosity is not yet used, the improved sampling can be seen to lead to a better convergence for both residual norms. However, at later stages the averaging-of-apparent-viscosity leads to ever decreasing effects of the improved sampling. After 20 Newton steps of the coupled problem, the difference in convergence levels obtained by adding more realization or extending the sampling time can be seen to be within the range of scatter for the different realizations of the coupled problem.

Increasing the number of realizations for each micro-scale MD problem and/or increasing the sampling time can be expected to reduce the statistical scatter in the predicted shear stresses for an imposed velocity gradient. This effect is demonstrated for the predicted apparent viscosity in Figures 12, 13. After approximately five Newton relaxation steps, only very small changes in the

DOI: 10.1002/fld

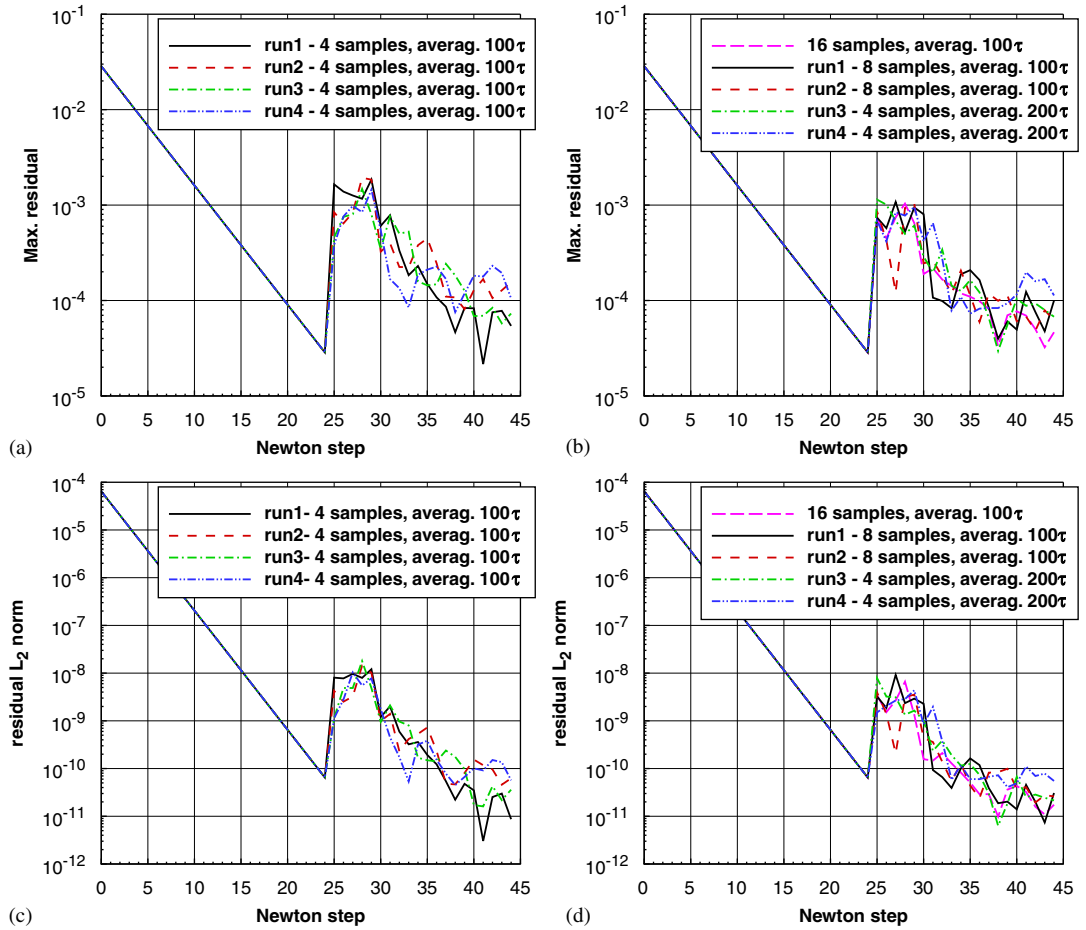

Figure 11. Convergence of 3D channel flow simulation. Finite-volume discretization with 40×40 cells, under-relaxed Newton relaxation ($\omega = 0.25$): (a) maximum residual, baseline sampling; (b) maximum residual, improved sampling; (c) L_2 norm, baseline sampling; and (d) L_2 norm, improved sampling.

Table I. Velocity gradients in cell face centres for which micro-scale problems are defined.

Cell face	0	1	2	3	4	5
du/dy $[\sigma^{-1}]$	0.007938	0.020028	0.005757	0.016619	0.004710	0.014220
du/dz $[\sigma^{-1}]$	0.000473	0.0	0.012865	0.010309	0.024205	0.020433

cell-face velocity gradients occur. Therefore, the figures also show the level of convergence that can be obtained for the apparent viscosity for a prescribed velocity gradient as function of the employed sampling technique. The averaging-of-apparent viscosity can be clearly seen to improve the level of convergence. Results are shown for four of the six MD micro-problems, which have imposed

DOI: 10.1002/fld

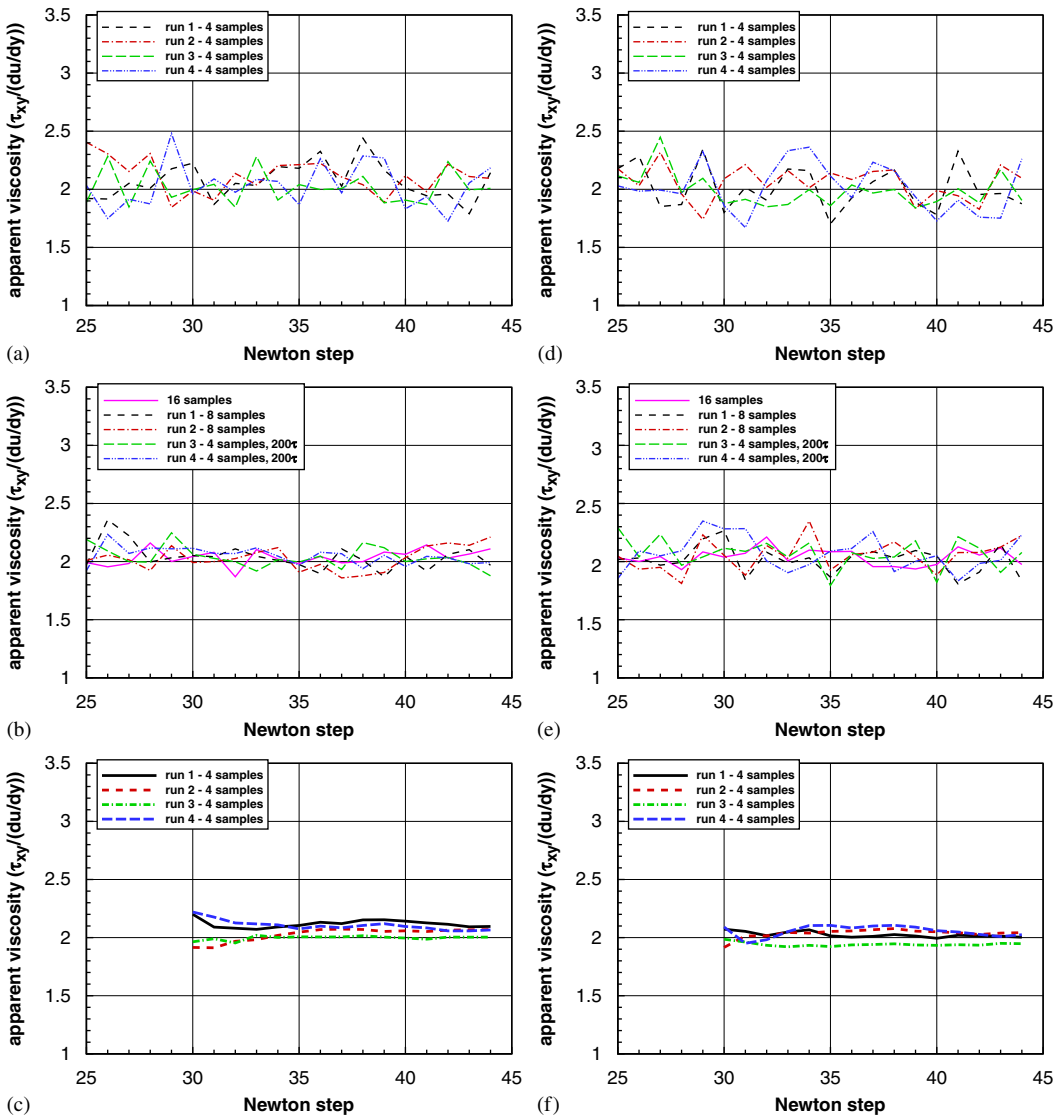

Figure 12. Evolution of apparent viscosity sampled from MD results in simulation of 3D channel flow. (a–c) cell face 1 and (d–f) cell face 3. Finite-volume discretization on stretched Cartesian mesh with 40×40 cells in cross-flow direction, under-relaxed Newton relaxation ($\omega = 0.25$).

velocity gradients of decreasing magnitude. This is reflected in the larger scatter in the predicted apparent viscosity for the second column relative to the first in both figures as well as the higher levels of scatter in Figure 12 relative to those in Figure 13. Results are compared for the simulations using the 'baseline' sampling (top row of figures) and for the simulations with either more MD realizations or extended sampling duration (middle row of figures). These predicted values of the apparent viscosity are compared with the one resulting from the averaging-of-apparent-viscosity

DOI: 10.1002/fld

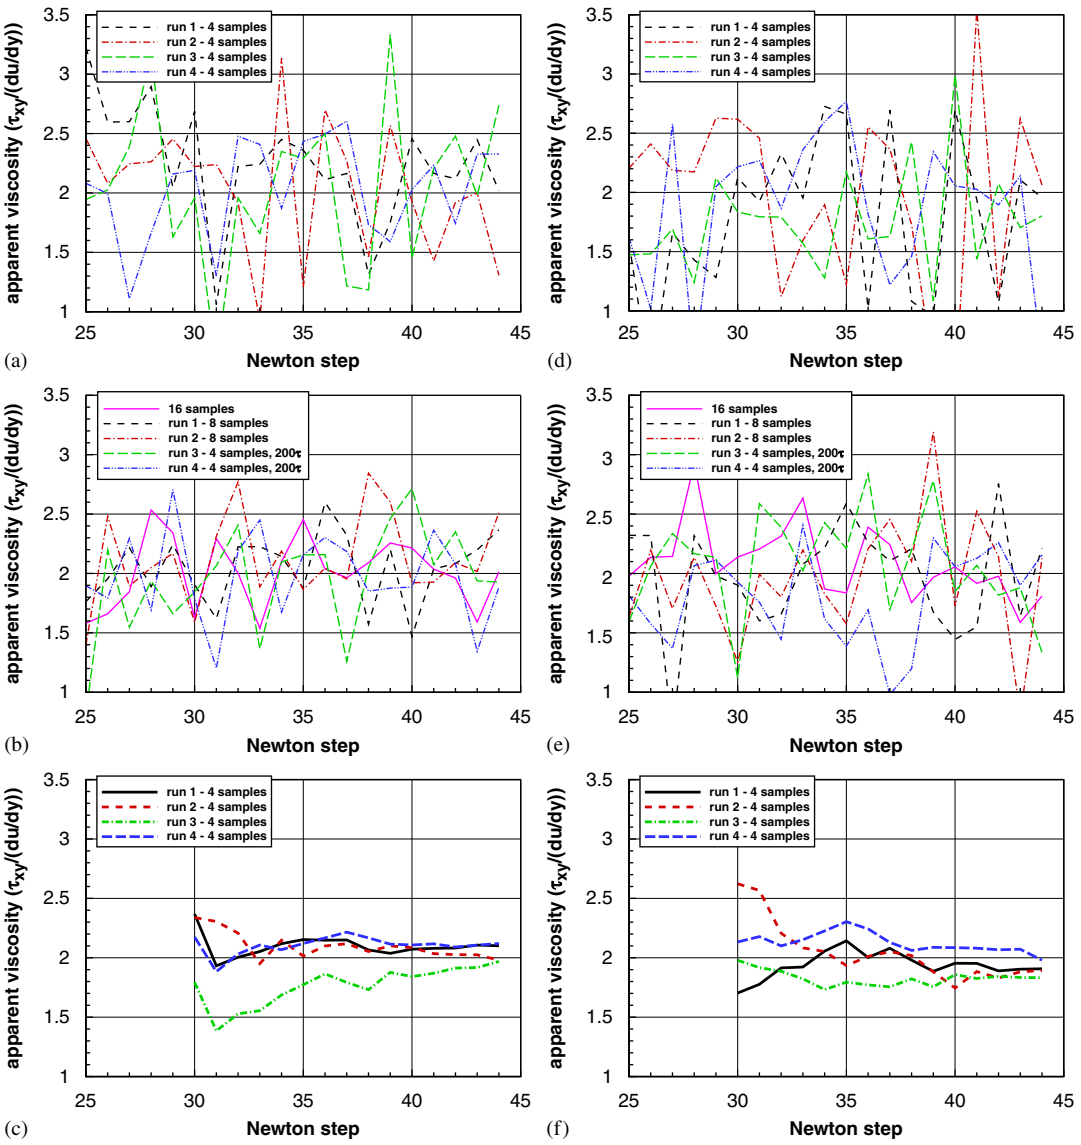

Figure 13. Evolution of apparent viscosity sampled from MD results in simulation of 3D channel flow. (a–c) cell face 2 and (d–f) cell face 4. Finite-volume discretization on stretched Cartesian mesh with 40×40 cells in cross-flow direction, under-relaxed Newton relaxation ($\omega = 0.25$).

approach (shown in the bottom row). Obviously, the improved sampling significantly reduced the scatter in the predicted apparent viscosity. However, the reduction of the scatter obtained by the averaging-of-apparent-viscosity approach outweighs this improvement, particularly at later stages in the simulation.

DOI: 10.1002/fld

7. DISCUSSION

In this section, the three-dimensional fully developed channel flow considered in the previous section is analysed using MD. The aim is to analyse the cost of the developed hybrid Navier–Stokes/MD method relative to a full MD simulation. For flows at the nano-scale, full MD solutions are affordable, and have been presented in the last two decades [3–6]. Clearly, for increasing length scales, this type of modelling becomes excessively expensive. The relative cost of full MD compared with the proposed hybrid method as function of the considered flow length scale is analysed in this section. In addition, the parallel efficiency of the developed framework is assessed for the considered cases.

7.1. Full MD simulation of channel flow

The fully developed channel flow is simulated here using a periodic Poiseuille problem [33]. In this approach, the mean flow is established using a volume source term, which replaces the pressure gradient driving the equivalent continuum channel flow. Figure 14 shows the considered geometry, with the mean flow in the homogeneous x-direction. In the y-, z-plane, the domain is divided in four sub-domains, while periodic boundary conditions are imposed in all three directions. The volume terms have opposite signs in neighbouring sub-domains, which leads to the time-averaged flow shown in the figure. Each of the four sub-domains can be considered as the equivalent of a fully developed channel flow. On average the velocity on the boundaries of the sub-domains will be zero; however, it should be noted that fluid–solid interactions are not modelled, which is consistent with the hybrid Navier–Stokes/MD simulations. For MD simulations of the channel flow, cubic domains of two different sizes are considered, the smaller with length 40σ in each direction and the larger one with edge length 80σ. For both domains, the macroscopic quantities are obtained by defining a rectangular uniform mesh of particle bins, with 80×80 bins in the cross-flow direction. Using the periodic Poiseuille approach, a macroscopic velocity field on a uniform 40×40 mesh can thus be obtained for the equivalent channel flow domain. For the smaller domain, four realizations of the MD simulation were obtained, while two realizations were obtained for the larger domain.

7.2. Analysis of computational cost

As a first step in comparing the cost of the developed hybrid Navier–Stokes/MD method for the considered test cases with that for a simulation using MD in the entire domain, the uncertainty in the mean velocity field predicted by the full MD simulation is analysed. The uncertainty in the obtained velocity field is assessed here by computing the predicted cross-flow velocity components as function of the sample duration as well as the number of independent realizations. The 'scatter' in the v and w velocity components, which should be zero, is shown in Figure 14(c–d) for the two domains. Three curves are shown in the plots. The first indicates the maximum value of v and w velocity components in all particle bins as function of sample duration. The second results shows the maximum value of the velocity components after ensemble averaging of the results in the homogeneous direction as well as over the independent realizations. Finally, the periodic spatial nature of the problems is used to combine the results for each of the four parts of the domain with opposite source sign. For the smaller domain, the scatter in cross-flow velocity (scaled with centre-of-channel velocity) is approximately 1.0% after sampling over 400τ, while for the larger

DOI: 10.1002/fld

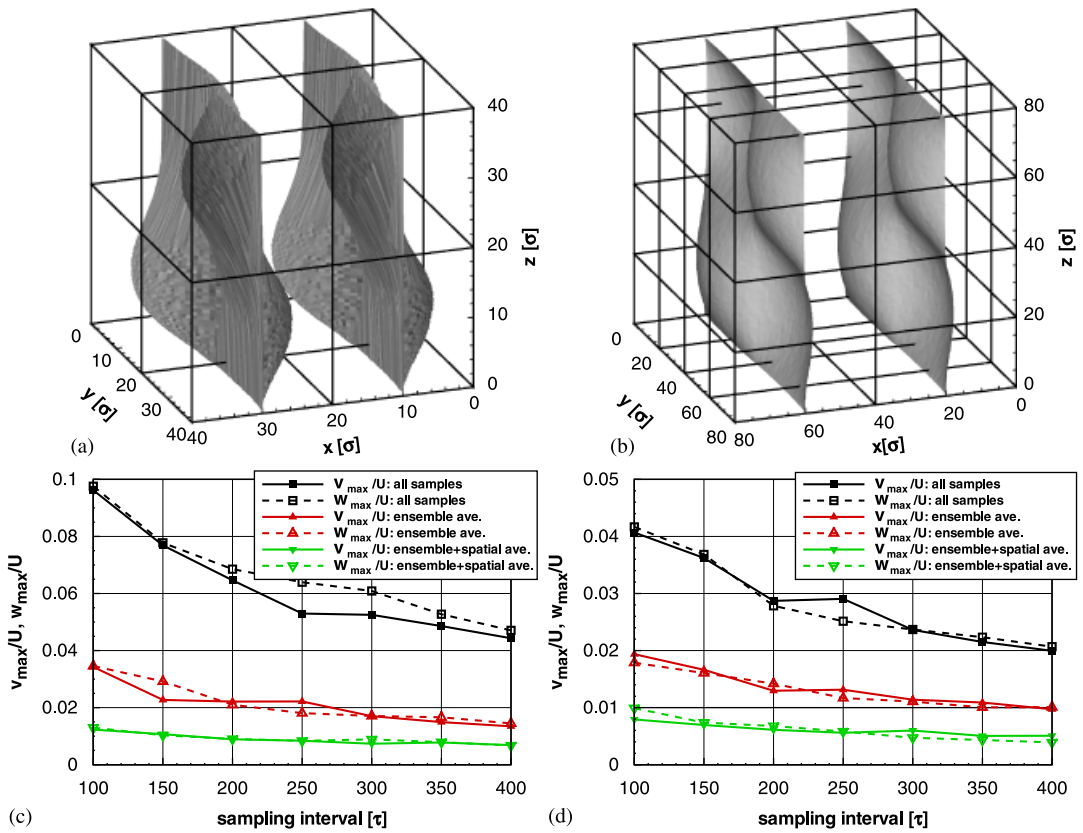

Figure 14. MD simulation of periodic Poiseuille problem. The mean flow is established using volume source terms replacing pressure gradient used in the continuum flow simulation. The domain is divided in four periodic rectangular regions (i.e. 2×2 in the y-, z-plane), with the source term changing sign across the periodic boundaries. This ensures a zero mean velocity in the x-direction for the full domain: (a) $40\times40\times40$ domain $-51\,200$ particles; (b) $80\times80\times80$ domain $-409\,600$ particles; (c) v, w scatter, $40\times40\times40$ domain; and (d) v, w scatter, $80\times80\times80$ domain.

domain, this is reduced to around 0.5% due to the increased number of particles which on average reside in each particle bin.

Equation (9) for the residual in the macro-scale finite-volume discretization can be used to estimate the maximum residual that corresponds to a given uncertainty in the velocity field. Assuming that the maximum residual occurs at the domain boundaries, where halo cell values are used to impose no-slip conditions, a change of 0.5% in the u-velocity component results in a residual increase due to the flux on the wall cell face of approximately $1.5\cdot10^{-3}$ based on the used cell size, Newtonian viscosity and the velocity in the channel centre. The other cell faces will have additional, but smaller, contributions. Hence, this conservative estimate shows that a maximum residual of $O(10^{-3})$ in Figures 10, 11 corresponds to an uncertainty in the solution of the velocity significantly smaller than 0.5%. Therefore after about five Newton steps in the coupled system, the uncertainty in velocity in the hybrid Navier–Stokes/MD solution is significantly smaller than

DOI: 10.1002/fld

in the full MD solution after 400τ sampling. The ratio of computational overhead for the hybrid Navier–Stokes/MD and the full MD simulations can be estimated as

$$\frac{\text{cost}_{\text{hybrid}}}{\text{cost}_{\text{MD}}} \approx \frac{n_{\text{Newton}} \times n_{\text{micro}} \times n_{\text{real,micro}} \times (n_{\text{equi}} + n_{\text{samp,micro}}) \times L_{\text{micro}}^3}{n_{\text{real,MD}} \times (n_{\text{equi}} + n_{\text{samp,MD}}) \times L_{\text{MD}}^3} \tag{10}$$

where L_{micro} and L_{MD} are the domain sizes of the micro-scale MD domain and full MD domain, respectively. n_{micro} is the number of micro-scale problems, $n_{\text{real,micro}}$ and $n_{\text{real,MD}}$ are the number of realizations of the micro-scale and full-scale MD simulations, respectively. n_{equi} is the number of time steps for equilibration prior to the sampling phase of the MD simulations, which uses $n_{\text{samp,micro}}$ and $n_{\text{samp,MD}}$ steps, respectively, for the micro-scale and full-scale MD simulations. For the coupled Navier–Stokes/MD, the parameters for the 'baseline' sampling are assumed, i.e. $n_{\text{samp,micro}} = 150\,000$, while $n_{\text{samp,MD}} = 450\,000$. The cost ratio then becomes, approximately, 0.27 and 0.07, for the smaller and larger full MD domains, respectively. For increasing domains this ratio will decrease further, when the assumption is made that only a small fraction of the cell faces in the macro-scale discretization requires viscous fluxes evaluated through MD.

7.3. Parallel execution

The channel flow simulations using the proposed hybrid method discussed in Section 6 were conducted in parallel on a Linux cluster consisting of 2.4 GHz Pentium processors connected by a Gigabit network. Each of the computed cases comprises a single finite-volume macro-scale problem and either 24, 48 or 96 micro-scale MD problems. For each Newton relaxation step, the CPU time required for an MD micro-scale problem exceeds that for the macro-scale finite-volume problem by at least three orders of magnitude. In the present parallel execution set-up, the micro-scale problems were distributed over either 12 or 24 processors, while the macro-scale problem was simply replicated on each processor for simplicity. Hence, each processors computed 1, 2 or 4 micro-scale problems per Newton relaxation step. Therefore, in this set-up, each of the micro-scale MD problems is computed on a single processor. The communication between the processors was restricted to collective operations needed in exchanging and averaging of shear stresses for the micro-scale problems predicted on different processors. Owing to the negligible overhead in the communication as well as the large ratio of micro-scale and macro-scale CPU times, the chosen execution method is practically embarrassingly parallel. The parallel solution strategy for the hybrid Navier–Stokes/MD method is therefore very well suited for the increasingly important *grid-computing* approach, in which computations are performed on a multitude of multi-processor computers that are connected by a relatively slow network, compared with the speed of the connection between the processors within the multi-processor computers.

The parallel efficiency of the MD method developed within the proposed framework was analysed for channel flow simulations using the periodic Poiseuille method, as discussed in Section 7. The dependence of the parallel efficiency on the chosen domain size as well as the used cut-off length in the cell linked-list algorithm were investigated. For these tests, a Linux cluster consisting of 2.4 GHz Opterons connected by a Gigabit network was used. Table II presents the parallel speed-up for up to 32 processors relative to serial execution. Cubic domains of two sizes are considered, the smaller one with length 80σ is each direction, while the larger one has domain edges of length 160σ. For the larger domain, the effect of the selected cut-off length in the linked-list algorithm is investigated. For the smaller domain considered, good parallel efficiency can be observed for up to eight processors, while the larger domain scales well up to 16 processors. The increase of

DOI: 10.1002/fld

Table II. Parallel speed-up for MD simulation of channel flow (periodic Poiseuille).

Domain	r_c	Link-cell mesh	4 proc.	8 proc.	16 proc.	32 proc.
$80^3[\sigma^3]$	4.0σ	20^3	3.58	6.60	11.0	20.4
$160^3[\sigma^3]$	2.5σ	64^3	3.56	6.65	11.9	21.8
$160^3[\sigma^3]$	4.0σ	40^3	3.81	7.38	14.2	27.0

the cut-off length in the cell linked-list algorithm significantly increases the volume of data that need to be communicated between the processors. However, the additional computational cost in the inter-particle potential computation (changing r_c from 2.5 to 4 LJ units more than quadruples the cost of this step) results in an increased speed-up for the larger cut-off length.

8. CONCLUSIONS

The development of a computational framework for multi-physics flow simulations has been described, along with a discussion of two example applications. In these examples, micro-scale MD problems are constructed to provide the viscous stresses for the Navier–Stokes solutions. The first problem was the planar Poiseuille flow, for which the restriction to Newtonian flow has been eliminated and the viscous fluxes on each cell face in the finite-volume discretization are evaluated using MD. The second example dealt with fully developed three-dimensional channel flow, with molecular level modelling of the shear stresses in a group of cells in the domain corners. The direct use of shear stresses evaluated with MD simulations leads to an unacceptable level of convergence of the residuals in the finite-volume method. This lack of convergence can be attributed to the scatter in the MD stresses due to the sampling of a finite ensemble over a limited interval. A number of solutions to this problem have been discussed in the present work. The averaging of the apparent viscosity for each cell face, i.e. the ratio of the shear stress predicted from MD and the imposed velocity gradient, over a number of Newton relaxation steps has been shown to be a simple but effective method to reach a good level of convergence of the coupled system. In the future, efforts are directed towards applications of the method to more complex flows and the assessment of the framework for generic, multi-physics problems.

REFERENCES

1. Allen M, Tildesly D. *Computer Simulation of Liquids*. Clarendon Press: Oxford, 1987.
2. Frenkel D, Smit B. *Understanding Molecular Simulation*: *From Algorithms to Applications*. Academic Press: San Diego, 2001.
3. Koplik J, Banavar J. Continuum deductions from molecular hydrodynamics. *Annual Review of Fluid Mechanics* 1995; **27**:257–292.
4. Koplik J, Banavar J. Corner flow in the sliding plate problem. *Physics of Fluids* 1995; **7**:3118–3125.
5. Din X, Michaelides E. Kinetic theory and molecular dynamics simulations of microscopic flows. *Physics of Fluids* 1997; **9**(12):3914–3925.
6. Soong C, Wang S, Tzeng P. Molecular dynamics simulation of rotating fluids in a cylindrical container. *Physics of Fluids* 2004; **16**(8):2814–2827.
7. Zhang J, Todd B, Travis K. Viscosity of confined inhomogeneous nonequilibrium fluids. *Journal of Chemical Physics* 2004; **121**:10778–10786.

DOI: 10.1002/fld

8. Ren W, Weinan E. Heterogeneous multiscale method for the modeling of complex fluids and micro-fluidics. *Journal of Computational Physics* 2005; **204**:1–26.
9. Lennard-Jones J. Cohesion. *Proceedings of the Physical Society* 1931; **43**:461–482.
10. Qian T, Wang X. Driven cavity flow: from molecular dynamics to continuum hydrodynamics. *SIAM Multiscale Modeling and Simulation* 2005; **3**(4):749–763.
11. Nie X, Chen S, Weinan E, Robbins M. A continuum and molecular dynamics hybrid method for micro- and nano-fluid flow. *Journal of Fluid Mechanics* 2004; **500**:55–64.
12. Nie X, Chen S, Robbins M. Hybrid continuum-atomistic simulation of singular corner flow. *Physics of Fluids* 2004; **16**(10):3580–3591.
13. Steijl R, Barakos G. Parallelisation of CFD methods for multi-physics problems. *European Conference on Computational Fluid Dynamics* (*ECCOMAS CFD 2006*), Egmond aan Zee, The Netherlands, September 2006.
14. Steijl R, Barakos G. Development of parallel computational framework for multi-physics flow simulations. *Fifth European Congress on Computational Methods in Applied Sciences and Engineering* (*ECCOMAS 2008*), Venice, Italy, 30 June–5 July 2008.
15. Plimpton S. Fast parallel algorithms for short-range molecular dynamics. *Journal of Computational Physics* 1995; **117**:1–19.
16. Grama A, Kumar V, Sameh A. Scalable parallel formulations of the Barnes–Hut method for n-body simulations. *Parallel Computing* 1998; **24**:797–822.
17. Kalé L, Skeel R, Bhandarkar M, Brunner R, Gursoy A, Krawetz N, Phillips J, Shinozaki A, Varadrajan K, Schulten K. NAMD2: greater scalability for parallel molecular dynamics. *Journal of Computational Physics* 1999; **151**:283–312.
18. Hariharan B, Aluru S. Efficient parallel algorithms and software for compressed octrees with applications to hierarchical methods. *Parallel Computing* 2005; **31**:311–331.
19. Matthey T, Cickovski T, Hampton S, Ko A, Qun M, Nyerges M, Raeder T, Slabach T, Izaguirre J. ProtoMol: an object-oriented framework for prototyping novel algorithms for molecular dynamics. *ACM Transactions on Mathematical Software* 2004; **30**:237–265.
20. Hess B, Kutzner C, van der Spoel D, Lindahl E. GROMACS 4.0: algorithms for highly efficient load-balanced, and scalable molecular simulation. *Journal of Chemical Theory and Computation* 2008; **4**:435.
21. Huber G, McCammon J. OOMPAA—object-oriented model for probing assemblages of atoms. *Journal of Computational Physics* 1999; **151**:264–282.
22. Malan A, Lewis R. On the development of high-performance C++ object-oriented code with application to an explicit edge-based fluid dynamics scheme. *Computers and Fluids* 2004; **33**:1291–1304.
23. Hatakeyama M, Watanabe M, Suzuki T. Object-oriented fluid flow simulation system. *Computers and Fluids* 1998; **27**:581–597.
24. Koumoutsakos P. Multiscale flow simulations using particles. *Annual Review of Fluid Mechanics* 2005; **37**: 457–487.
25. Verlet L. Computer experiments of classical fluids. I. Thermodynamical properties of Lennard-Jones molecules. *Physical Review* 1967; **159**:98.
26. Lees A, Edwards S. The computer study of transport processes under extreme conditions. *Journal of Physics C* 1972; **5**:1921–1929.
27. Thomas J, Rowley R. Transient molecular dynamics simulations of viscosity for simple fluids. *Journal of Chemical Physics* 2007; **127**:174510.
28. Berendsen H, Postma J, van Gunsteren W, Dinola A, Haak J. Molecular-dynamics with coupling to an external bath. *Journal of Chemical Physics* 1984; **81**:3684–3690.
29. Irving J, Kirkwood J. The statistical mechanical theory of transport processes IV. *Journal of Chemical Physics* 1950; **18**:817–829.
30. Hess B. Determining the shear viscosity of model liquids from molecular dynamics simulations. *Journal of Chemical Physics* 2002; **116**:209–217.
31. Meier K. Transport coefficients of the Lennard-Jones model Fluid. I. Viscosity. *Journal of Chemical Physics* 2004; **121**:3671–3687.
32. Viscardy S, Servantie J, Gaspard P. Transport and Helfand moments in the Lennard-Jones fluid. I. Shear viscosity. *Journal of Chemical Physics* 2007; **126**:184512.
33. Backer J, Lowe C, Hoefsloot H, Iedema P. Poiseuille flow to measure the viscosity of particle model fluids. *Journal of Chemical Physics* 2005; **122**:154503.

DOI: 10.1002/fld