Multiscale simulation of nanofluidic networks of arbitrary complexity

David Stephenson · Duncan A. Lockerby ·
Matthew K. Borg · Jason M. Reese

Received: 1 April 2014 / Accepted: 26 August 2014 / Published online: 12 September 2014
© The Author(s) 2014. This article is published with open access at Springerlink.com

Abstract We present a hybrid molecular-continuum method for the simulation of general nanofluidic networks of long and narrow channels. This builds on the multiscale method of Borg et al. (Microfluid Nanofluid 15(4):541–557, 2013; J Comput Phys 233:400–413, 2013) for systems with a high aspect ratio in three main ways: (a) the method has been generalised to accurately model any nanofluidic network of connected channels, regardless of size or complexity; (b) a versatile density correction procedure enables the modelling of compressible fluids; (c) the method can be utilised as a design tool by applying mass-flow-rate boundary conditions (and then inlet/outlet pressures are the output of the simulation). The method decomposes the network into smaller components that are simulated individually using, in the cases in this paper, molecular dynamics micro-elements that are linked together by simple mass conservation and pressure continuity relations. Computational savings are primarily achieved by exploiting length-scale separation, i.e. modelling long channels

as hydrodynamically equivalent shorter channel sections. In addition, these small micro-elements reach steady state much quicker than a full simulation of the network does. We test our multiscale method on several steady, isothermal network flow cases and show that it converges quickly (within three iterations) to good agreement with full molecular simulations of the same cases.

Keywords Multiscale simulations · Hybrid methods · Molecular dynamics · Coupled solvers · Scale separation · Nanofluidics

1 Introduction

There are a number of engineering applications of fluid dynamics that require modelling with nanoscale or micro-scale resolution, for example, lab-on-a-chip devices (Jiang et al. 2011), nanotube membranes for seawater desalination (Mattia and Gogotsi 2008) or air purification (Mantzalis et al. 2011), and miniaturised heat exchangers for cooling electronic circuits (Yarin et al. 2009). In any simulation, the smallest dimension of a problem will determine the resolution required, and these applications are all multiscale: the length scale in one direction (usually the streamwise direction) is orders of magnitude greater than in the others. For example, the length of a nanotube in a desalination membrane (~ 2 mm) can be far larger than its diameter (~ 2 nm). This makes most micro solvers, such as molecular dynamics (MD), very computationally inefficient because the computation time grows dramatically with the size of the domain. The use of continuum fluid hydrodynamics, such as the Navier–Stokes–Fourier (NSF) equations, is also inadequate for these types of problems. While the NSF equations have been shown to still be legitimate

D. Stephenson · D. A. Lockerby (✉)
School of Engineering, University of Warwick,
Coventry CV4 7AL, UK
e-mail: D.Lockerby@warwick.ac.uk;
duncan.lockerby@warwick.ac.uk

D. Stephenson
e-mail: david.stephenson@warwick.ac.uk

M. K. Borg
Department of Mechanical and Aerospace Engineering,
University of Strathclyde, Glasgow G1 1XJ, UK
e-mail: matthew.borg@strath.ac.uk

J. M. Reese
School of Engineering, University of Edinburgh,
Edinburgh EH9 3JL, UK
e-mail: jason.reese@ed.ac.uk

Fig. 1 Schematics of different types of hybrid multiscale methods applied to (a) a generic internal-flow problem: (b) the domain decomposition method (DD), (c) the heterogeneous multiscale method

(HMM), and (d) the internal-flow multiscale method (IMM). Reproduced from Borg et al. (2013a)

for Newtonian flows at the nanoscale, they break down for non-Newtonian flows (with unknown constitutive relations), flows dominated by surface effects (e.g. molecular layering or large slip), or flows that exhibit thermodynamic non-equilibrium due to high gradient processes in the bulk fluid. A number of so-called hybrid multiscale methods have therefore been developed that aim to combine the best features of both micro and macroscale solvers: these strike a compromise between the accuracy of MD (in places where NSF is invalid) and the computational speed of NSF (where MD would be too computationally expensive).

The three main types of hybrid method are domain decomposition (O'Connell and Thompson 1995; Hadjiconstantinou and Patera 1997; Wagner et al. 2002; Delgado-Buscalioni and Coveney 2003a; Nie et al. 2004), the heterogeneous multiscale method (Ren and E 2005; Yasuda and Yamamoto 2008; E et al. 2009; Asproulis et al. 2012; Alexiadis et al. 2013), and the internal-flow multiscale method (Borg et al. 2013a, b; Patronis et al. 2013). In domain decomposition (DD), the flow domain is spatially divided so that there are micro solvers at the surfaces and neighbouring fluid regions, while a macro solver is applied for the bulk fluid (see Fig. 1b). Fluid state properties or fluxes are exchanged in an overlap region that couples the micro

and macroscale solvers to each other. The DD method is mostly used for flows that exhibit near-surface phenomena. Its disadvantage is that for long, narrow channels there is no great computational saving because the entire length of the bounding wall must be simulated using the micro solver, and the bulk region is relatively small compared to what can be considered the near-wall region. This restricts the dimensions of systems that DD can be useful for. In addition, there is the complicated issue of the requirement for non-periodic boundary conditions (NPBCs) for the micro solver at the overlap region, and these may not be well established (see Mohamed and Mohamad (2010) and references therein). Furthermore, the accuracy of the coupling is significantly diminished if the fluid viscosities used in the micro and macro solvers are not exactly equal (Delgado-Buscalioni and Coveney 2003a).

In the heterogeneous multiscale method (HMM), the macro solver spans the entire domain, with micro solvers spatially dispersed to provide the micro resolution (see Fig. 1c). The macro solver is updated with local fluid information from the micro solvers. For example, at surfaces, slip velocity and shear stress information is passed to the macro solver, which negates the need for slip boundary conditions; in the bulk fluid, shear stress information

is provided, removing the need for presumed constitutive relations or transport coefficients (Ren and E 2005). In return, the micro solvers are constrained by the local strain rate (and the temperature for non-isothermal flows) measured from the macro solver. The main drawback of HMM is that for narrow flow channels, the micro solver regions would likely overlap, rendering the computational speed-up small and the accuracy lower than that of a pure micro solver. Problems with NPBCs remain and, for a fluid with unknown constitutive relations, the number of micro solvers to accurately represent the varying stress field would further reduce any computational speed-up.

The internal-flow multiscale method (IMM) (Borg et al. 2013b; Patronis et al. 2013) is a recent development to tackle the specific problem of simulating flow in long and narrow channels, for which (as outlined above) other hybrid techniques are less computationally efficient. The IMM hybridisation exploits length-scale separation between the gradually varying hydrodynamic processes occurring on scales of similar magnitude to the length of the channel, and the molecular processes occurring on scales of similar magnitude to the channel height or width. In IMM, a simple continuum description applies throughout the flow domain, but a number of short micro-solver elements are distributed along the length of the channel, each covering the entire channel height or width (see Fig. 1d). The number and position of micro-elements are chosen to be sufficient to resolve the geometrical features of the channel, but contain, in total, far fewer molecules or particles than a simulation of the entire domain. There is no direct communication between the micro and macro solvers. Instead, coupling is performed iteratively via simple mass conservation and pressure continuity over the entire channel: the micro solver provides a local mass flow rate and the macro-solver provides either a local pressure difference or a pressure gradient. This approach removes any need to pass surface information (such as slip) or constitutive information (such as viscosity) to the macro solver, as these are inherently modelled within the micro-element simulations; the mass flow rate is the only information that is needed by the macro formulation from the micro elements. Another benefit is that IMM enables periodic boundary conditions (PBCs) to be used in the micro-element simulations, and these are far more efficient and convenient. One current limitation of this method, though, is that it is only applicable to isothermal and steady-state flows, because the simulations require long averaging periods to obtain low-noise measurements prior to coupling.

Borg et al. (2013a) applied the IMM to the simulation of serial networks of high-aspect-ratio channels (here referred to as a Serial-Network IMM implementation, SeN-IMM). The SeN-IMM decomposes a network into *components*, of which there are two types: *channel components* (which

Fig. 2 Schematic of an example complex network with high-aspect-ratio channels that can be solved by our hybrid method proposed in this paper. The network is decomposed into components by adopting an IMM approach in the long channels and full molecular dynamics simulations of the junctions

have a high aspect ratio, and can thus be represented by a shorter, computationally cheaper, series of micro elements as illustrated in Fig. 1d) and *junction components* (of arbitrary geometry, such as a reservoir or channel defect, for which no scale separation can be exploited, and thus must be simulated in their entirety). The term *serial* refers to the fact that in Borg et al. (2013a), every component in the network could only have one inflow and one outflow, and so, the mass flow rate was constant in each micro element. A general network solver would require far greater practical applicability (e.g. for mixing channels and lab-on-a-chip devices) and is not trivial to implement. An example of a general network is shown in Fig. 2 and includes a bifurcating flow. Another limitation of the method proposed in Borg et al. (2013a) is that it is only appropriate for effectively incompressible fluid conditions. However, flows at the nanoscale display compressibility even at very low Mach numbers (Gad-El-Hak 2006).

In this paper, we therefore extend the previous SeN-IMM technique to encompass general network geometries (where each component can have any number of inlets and outlets) and incorporate an equation of state to account for fluid compressibility. We refer to this new method as the General Networks Internal-flow Multiscale Method (GeN-IMM). The other assumptions in Borg et al. (2013a) are maintained, i.e. isothermal and steady-state flows. Our new method can potentially be used in conjunction with any micro solver, but for this paper, we have chosen MD.

Another novel aspect of our method is that it has the capability to act as a design tool, rather than just a simulation tool. It enables mass flow rates (rather than just inlet/outlet pressures) to be specified as boundary conditions. This is important in applications where the mass flow rate must be controlled, but the required inlet or outlet pressure to generate it is not known, for example, in drug delivery.

The GeN-IMM provides computational savings in three ways: (1) long channels are replaced by hydrodynamically equivalent shorter channel micro elements; (2) these micro-elements reach steady state much faster than a simulation

Fig. 3 Schematics of (a) a simple Y-junction network with (b) the network decomposed into components; internal and external boundaries are highlighted

of the full network, due to their small size; and (3) the trial-and-error process necessary to determine pressure boundary conditions to achieve a desired mass flow rate is removed. Furthermore, as the micro elements are small and can be simulated independently of each other during each iteration, this new method is less dependent on access to supercomputer resources: micro-elements can be run on single graphics processing units (GPUs) instead of a large number of central processing units (CPUs). There is often a limit to the number of molecules able to be processed on a GPU, and so, large MD simulations cannot benefit from them.

2 The hybrid method for networks: GeN-IMM

We decompose a general fluidic network into *components*, with each being defined as either a *junction* component or a *channel* component (see Fig. 3 for an illustrative example of such a decomposition). Our main purpose is to enable the channel components to be simulated separately using short, periodic, and computationally cheaper MD simulations (this is the IMM approach). Each component has a number of ‘inlet/outlet’ *boundaries*; these are either *external boundaries* to the network or *internal boundaries* which connect neighbouring components (see Fig. 3).

To make the separate simulations of components consistent with conditions in the full network, the correct pressure values at all boundaries must be established. These must ensure that the mass flow rate ($\dot{m}$) and pressure (P) at all internal boundaries are consistent and that any external boundary conditions are satisfied, i.e.:

$$\dot{m}_{i,p} = -\dot{m}_{j,q}, \quad (1)$$

$$P_{i,p} = P_{j,q}, \quad (2)$$

where i and j refer to a pair of connecting internal boundaries, from neighbouring components p and q , respectively, and

$$\dot{m}_{i,q} = \dot{m}_B, \quad (3)$$

$$P_{i,q} = P_B, \quad (4)$$

where the subscript B denotes an external boundary condition, either on pressure or mass flow rate, on the i th boundary of the q th component. Note, by convention, the mass flow rate is treated as positive flowing out of the component, hence the minus sign in Eq. (1).

In steady-state flows, the net mass flow rate through all boundaries of one component is zero:

$$\sum_{i=1}^{W_q} \dot{m}_{i,q} = 0, \quad (5)$$

where W_q is the total number of boundaries of the q th component. This *local* mass conservation (5) is automatically guaranteed by any individual component simulation, but it is only in combination with (1) that *global* mass conservation is satisfied over the whole network.

The required pressure values are found iteratively, with successive estimations moving conditions closer to global mass conservation (within the constraints described above). To generate pressure values for the next iteration, a linear prediction of how mass flow rate will change in response to changes in pressure is used:

$$\begin{aligned} \dot{m}_{i,q} - \sum_{j=1(\neq i)}^{W_q} (P_{i,q} - P_{j,q}) K_{ij,q} \\ = \langle \dot{m}_{i,q} \rangle - \sum_{j=1(\neq i)}^{W_q} (\langle P_{i,q} \rangle - \langle P_{j,q} \rangle) K_{ij,q}, \end{aligned} \quad (6)$$

where the terms in angle brackets are measurements extracted from component MD simulations at the previous iteration; the terms $K_{ij,q}$ are flow-conductance coefficients between the boundaries i and j in component q (their calculation is discussed in Appendix 1). Equation (6), in combination with Eqs. (1)–(5), provides a system of linear equations with an equal number of unknowns. This system can be solved using a straightforward matrix inversion procedure (e.g. LU decomposition) to obtain values of pressure for the next iteration. Equation (6) ensures that as the component MD simulations approach global mass conservation (i.e. $\langle \dot{m}_{i,q} \rangle \rightarrow \dot{m}_{i,q}$), the values of pressure cease to be updated in subsequent iterations (i.e. $P_{i,q} \rightarrow \langle P_{i,q} \rangle$). Note,

this converged result is independent of the flow conductances ($K_{ij,q}$), which only affect the convergence rate and the approach to convergence.

Fluid compressibility in the network, due to large changes in pressure, can be allowed for in GeN-IMM by modifying the number of molecules in each component simulation according to an empirical equation of state relating average mass density to the boundary pressures of the component. The empirical equation is of the form:

$$\bar{\rho}_q = \sum_{b=0}^{\beta} a_{b,q} \bar{P}_q^b, \quad (7)$$

where $\bar{\rho}_q$ is the volume-averaged mass density in the q th component, $a_{b,q}$ are constants, β is the order of the polynomial equation, and

$$\bar{P}_q = \frac{1}{W_q} \sum_{i=1}^{W_q} P_{i,q}, \quad (8)$$

where $P_{i,q}$ is the pressure at the i th of W_q boundaries. The coefficients $a_{b,q}$ are determined from a presimulation of a periodic cube of fluid molecules. Pressure measurements for different densities are obtained using a standard Irving-Kirkwood calculation, and a β -order polynomial is then fitted to the data, which provides the coefficients. The values of $a_{b,q}$ for liquid argon can be found in Appendix 2; the same values are used in all the multiscale solutions of this paper. Equations (7) and (8) enable us to estimate the number of molecules that need to be added or removed (using the FADE algorithm of Borg et al. (2014), the USHER algorithm of Delgado-Buscalioni and Coveney (2003b), or the AdResS scheme of Praprotnik et al. (2005)) for the next iteration of each component simulation.

2.1 Simulating channel components

Solving Eqs. (1)–(6) provides the pressure values that satisfy external boundary conditions, that are consistent at internal boundaries, and that create a mass-flow response that is closer to being globally conservative than that of the previous iteration. We now discuss how these boundary conditions can be applied to the individual component simulations.

The channel components, by definition, are highly scale separated and can thus be simulated using shorter MD channels (which is the same as the IMM procedure of Borg et al. (2013b) and Patronis et al. (2013)). In the example of Fig. 3(b), the channel has a uniform cross section in the streamwise direction. Therefore, given the assumption that the pressure varies approximately linearly along the entire channel length L , it is sufficient to model the channel using a shorter periodic MD micro-element of length

Fig. 4 Schematic of the channel component from Fig. 3, demonstrating how a shorter micro MD simulation is used to represent the original channel

L' with a uniform cross-sectional area, as shown in Fig. 4. A weak streamwise pressure gradient is hydrodynamically equivalent to a constant streamwise body force because the momentum flux produced is identical. As such, the pressure drop across a long channel can be simulated by applying a uniform body force F_q to all fluid molecules in a periodic MD channel of arbitrary length, using a central-difference approximation:

$$F_q = \frac{(P_{i,q} - P_{j,q})m}{\bar{\rho}_q L}, \quad (9)$$

where m is the mass of a single molecule, and the direction of a positive force is from boundary i to boundary j . The computational savings accrued by using small MD micro-elements to model large-aspect-ratio channel components is the main source of speed-up in the GeN-IMM and is roughly proportional to L/L' for these components.

In certain liquid cases, and commonly in rarefied gases (Huang et al. 2007), the pressure distribution along long channels can be nonlinear. In these instances, the channel component is divided into multiple MD micro-elements, such that the combination of the linear pressure gradient over each micro-element is a good approximation of the overall pressure variation (see Borg et al. (2013b); Patronis et al. (2013)).

2.2 Simulating junction components

Junction components, unlike channels, cannot be represented using smaller MD simulations; there is no obvious scale separation that can be exploited, so the full component must be simulated.

However, what is not straightforward in MD is how to deal with the non-periodicity of most junction components

(e.g. the Y-junction in Fig. 3). There are no available MD ensembles to solve non-periodicity, even though a significant amount of progress has been made on developing artificial NPBCs (see Mohamed and Mohamad (2010) and references therein).

Here, we circumvent this problem by introducing an artificial region attached to the main component, which enables convenient PBCs to be used, as shown in Fig. 5. The correct pressure boundary conditions for the real region are established by applying body forces in localised regions of constant cross section within the artificial region, which create sharp jumps in pressure, $\Phi_{i,q}$ (and introduces more unknowns into the system of equations). As in Eq. (6), these body forces are determined iteratively by making a linear prediction of how the mass flow rate and pressure in the real domain will be affected by the body-force-generated pressure jumps $\Phi_{i,q}$:

$$\dot{m}_{i,q} + \sum_{j=1(\neq i)}^{W_q} \Delta P_{ij,q} L_{ij,q} = \langle \dot{m}_{i,q} \rangle + \sum_{j=1(\neq i)}^{W_q} \langle \Delta P_{ij,q} \rangle L_{ij,q}, \quad (10)$$

where $\Delta P_{ij,q} = (P_{i,q} - \Phi_{i,q}) - (P_{j,q} - \Phi_{j,q})$, $L_{ij,q}$ are flow-conductance coefficients, and the angular brackets denote a measurement extracted from the MD simulation at the previous iteration. The exact relationship between the body forces $F_{i,q}$ and pressure jumps $\Phi_{i,q}$ is given in Appendix 3. The sign differences between Eqs. (6) and (10) arise because the mass flow rate at a boundary is deemed to be positive if it flows out of the component.

At one boundary in each component (chosen arbitrarily), no body force is applied, i.e. $\Phi_{1,q} = 0$, so as not to over-constrain the simulation.

2.3 Algorithm

We now outline the iterative algorithm for general multi-scale network configurations:

1. Approximate the flow-conductance coefficients $K_{ij,q}$, $L_{ij,q}$ for each component. The terms in angular brackets in Eqs. (6) and (10) are assumed to be zero.
2. Solve the set of linear Eqs. (1)–(6) and (10), using matrix inversion, for the predicted mass flow rates $\dot{m}_{i,q}$, pressures $P_{i,q}$ and junction pressure jumps $\Phi_{i,q}$.
3. Solve Eqs. (7)–(9) for the mass densities ρ_q and channel component body forces F_q , using the predicted values of pressure $P_{i,q}$ previously calculated.
4. Run all micro-element MD simulations with the new body forces and updated average densities until steady state. At steady state, measure time-averaged mass flow rates $\langle \dot{m}_{i,q} \rangle$ and pressures $\langle P_{i,q} \rangle$ at every boundary (for junctions), and the mean pressure $\langle \bar{P}_q \rangle$ (for channels).

Fig. 5 Schematics of (a) the junction component from Fig. 3 with body forces $F_{i,q}$ and $F_{j,q}$ applied in the artificial region and (b) the periodic MD micro-element simulation setup of the same component

These measured properties are used in Eqs. (6) and (10) for the next iteration. Note, for channel components, the term $(P_{i,q} - P_{j,q})$ is obtained directly from Eq. (9).

5. Update the flow-conductance coefficients $K_{ij,q}$ and $L_{ij,q}$.
6. Repeat from step 2 until a convergence criterion is met for the mass flow rate at a given boundary:

$$\zeta_{i,q} < \zeta^{\text{tol}}, \quad \text{with} \quad \zeta_{i,q} = \left| \frac{[\dot{m}_{i,q}]_l - [\dot{m}_{i,q}]_{l-1}}{[\dot{m}_{i,q}]_l} \right|, \quad (11)$$

or for the pressure:

$$\zeta_{i,q} < \zeta^{\text{tol}}, \quad \text{with} \quad \zeta_{i,q} = \left| \frac{[P_{i,q}]_l - [P_{i,q}]_{l-1}}{[P_{i,q}]_l} \right|, \quad (12)$$

depending on what the boundary constraints are. Here ζ^{tol} is a predetermined convergence tolerance and the superscripts l and $l-1$ denote values calculated at the current and the previous iteration, respectively.

2.4 Molecular dynamics

In order to validate our hybrid technique, we solve full networks with non-equilibrium MD simulations (Rapaport 2004; Allen and Tildesley 1987) using the mdFoam solver (Macpherson and Reese 2008; Borg et al. 2010) developed in OpenFOAM, an open-source set of C++ libraries for solving sets of differential equations in parallel (www.openfoam.org). We use the same MD solver in the iterative scheme outlined above, applying body force and density constraints to the micro-element simulations.

For simplicity, and in order to demonstrate our technique, we simulate spherically symmetric monatomic particles (hereon referred to as ‘molecules’) of liquid argon, which interact through Lennard-Jones pairwise potentials.

However, our method is generally applicable to more complex and realistic molecular models. Solid walls in all the MD simulations consist of rigid molecules of argon that are kept fixed in space and time. The dynamics of the liquid argon molecules are described by Newton's equations of motion, with an added external force, $\mathbf{f}^{(ext)}$:

$$\frac{d}{dt}\mathbf{r}_k = \mathbf{v}_k, \quad (13)$$

$$m_k \frac{d}{dt}\mathbf{v}_k = \mathbf{f}'_k + \mathbf{f}^{(ext)}, \quad (14)$$

where k is a molecule in the system and $\mathbf{r}_k = (x_k, y_k, z_k)$ is its time-instantaneous position in a fixed Cartesian coordinate reference system. Each molecule has a mass m_k and a velocity $\mathbf{v}_k = (u_k, v_k, w_k)$. The total intermolecular force on each molecule $\mathbf{f}'_k = -\sum_j \nabla U(r_{kj})$ is determined from neighbouring molecules j , where $U(r_{kj})$ is the pair potential and $r_{kj} = |\mathbf{r}_k - \mathbf{r}_j|$ is the pair-molecule separation. Equations (13) and (14) are integrated in space and time using the Velocity Verlet method with a time-step of 5.4 fs. We adopt the Lennard-Jones (LJ) 12–6 potential, widely used to model simple liquids,

$$U_{12-6}(r_{kj}) = \begin{cases} 4\epsilon \left[\left(\frac{\sigma}{r_{kj}} \right)^{12} - \left(\frac{\sigma}{r_{kj}} \right)^6 \right] & \text{if } r_{kj} < r_c \\ 0 & \text{if } r_{kj} \geq r_c, \end{cases} \quad (15)$$

where σ and ϵ are the length and energy characteristics of the potential, and $r_c = 1.36$ nm is the cut-off separation. The σ and ϵ properties for the liquid–liquid (l–l) and wall–liquid (w–l) interactions are taken from Thompson and Troian (1997), with the intention to generate slip at solid–liquid interfaces. The values for these are, $\sigma_{l-l} = 3.4 \times 10^{-10}$ m, $\epsilon_{l-l} = 1.65678 \times 10^{-21}$ J, and $\sigma_{w-l} = 2.55 \times 10^{-10}$ m, $\epsilon_{w-l} = 0.33 \times 10^{-21}$ J. The mass density of the wall molecules is $\rho_w = 6.809 \times 10^3$ kg/m³, and the mass of one molecule is 6.6904×10^{-26} kg.

All our MD simulations are three-dimensional, with PBCs applied in every direction. The cases are all set up so that there are no gradients of properties in one direction perpendicular to the flow; so, the depth is chosen as a compromise between computational efficiency and the ability to generate sufficient data for averaging.

The external forcing $\mathbf{f}^{(ext)}$ in Eq. (14) is used both to generate pressure jumps (in the artificial regions of the junction components) and to represent linear pressure gradients (for channel components), as described previously. The heat generated indirectly by this forcing is removed to ensure a thermally homogeneous system; we utilise a Berendsen thermostat (Berendsen et al. 1984) to rescale the velocities of molecules. This thermostat operates on the thermal velocities, minimising the impact on the streamwise

velocity. The thermostat is implemented via localised bins over the entire MD domain in the streamwise and transverse flow directions, with bin sizes of 0.68 nm. Each bin has a target temperature of $T = 292.8$ K, using a time-constant $\tau_T = 21.61$ fs.

The net mass flow rate $\dot{m}_i$ (kg/s) is calculated by summing the number of molecules that cross a measurement plane in a given direction over a long averaging time Δt_{av} . For junction components, mass flow rate is measured at each boundary. The mass of the total number of molecules that cross the boundary travelling from the real region to the artificial region is counted as positive and that crossing in the opposite direction is counted as negative, i.e.

$$\langle \dot{m}_i \rangle = \frac{1}{\Delta t_{av}} \sum_k^{\delta N} m_k \text{sgn}(\mathbf{v}_k \cdot \hat{\mathbf{n}}_i), \quad (16)$$

where $\hat{\mathbf{n}}_i$ is the streamwise direction unit vector and δN is the total number of molecules that cross the boundary during the time period $t \rightarrow t + \Delta t_{av}$. The direction of crossing is obtained by the signum function, $\text{sgn}(x)$. For channel components, the mass flow rate is measured at the centre of the micro-element.

3 Results and discussion

The GeN-IMM method is tested on compressible pressure-driven flows through some simple network configurations. Due to the statistical noise created by thermal fluctuations, large pressure gradients are required in MD simulations of Poiseuille flows (Koplik et al. 1988; Travis et al. 1997) in order to achieve sufficiently low-variance data within a reasonable time frame.

Our results are validated via comparison to a full MD simulation of the entire network. In this way, any approximations made within the MD model are negated as they are the same in both the full and the hybrid cases. The pressure drops over the full MD networks are generated in the same manner as for junction micro-elements (i.e. the application of body forces in an artificial region). The networks are restricted to a relatively small size in order that these full MD simulations are not too computationally expensive. The full MD simulations are run in parallel on 48 CPUs, while the MD micro-elements for the hybrid solution are run on single GPUs.

The first network we analyse is a straight channel connecting two reservoirs i.e. a serial network, which is similar to a test case in Borg et al. (2013a); the second network is a bifurcating channel i.e. a general network, which demonstrates the novel capabilities of GeN-IMM.

Fig. 6 **a** Schematic of the straight channel multiscale network with **(b)** the hybrid MD GeN-IMM decomposition and **(c)** the full MD setup. Dimensions are in nanometres. The boundary numbers are labelled above each image

3.1 A straight channel geometry

The configuration of the first network is a relatively long, straight nanochannel connecting an inlet reservoir to an outlet reservoir, as shown in Fig. 6a. Figure 6b shows the hybrid decomposition, consisting of three micro-elements: an inlet reservoir, a short channel to represent the long channel, and an outlet reservoir. The extent of the entrance/exit channel sections of the reservoir micro-elements has been chosen conservatively to be roughly twice the channel flow entrance length from laminar macroscopic flow theory. This ensures the flow in micro-element #2 is fully developed and is not affected by expansion/contraction effects.

The length of the channel component is $L = 102$ nm, which is represented in the multiscale model by a micro-element of length $L' = 4.08$ nm, producing a length ratio of $L/L' = 25$. The height of the channel section is 3.4 nm, which is sufficiently small that there is non-continuum flow behaviour in the LJ liquid argon (e.g. molecular layering

and velocity slip at the liquid–solid interface) that would not be effectively captured by a standard Navier–Stokes fluid dynamics solution. The reservoir height (in the y -direction) is 6.8 nm, and the depth (in the z -direction) is a uniform 6.8 nm along the entire network and in each micro-element. A large pressure drop of ≈ 350 MPa is imposed over the network for the reason described above.

This network geometry is used to present two cases: case 1, where the multiscale method is constrained by external pressure boundary conditions at the inlet and outlet, i.e. boundaries #1 648 MPa, and #6 295 MPa, respectively (see Fig. 6a); case 2, where the external boundary conditions are an inlet pressure of 648 MPa and a mass flow rate 1.54 ng/s. This will demonstrate the versatility of the method as either a simulation tool (case 1) or a design tool (case 2).

The iterative algorithm previously outlined is performed for 5 iterations (although convergence occurs in fewer) to illustrate the numerical stability of the method. Figure 7 shows an example of how the measured mass flow rate and pressure develops with iteration number at boundary #1

Fig. 7 Straight channel case 1 multiscale measurements of mass flow rates and pressure at boundary #1, progressing with number of iterations. The pressure error bars are smaller than the size of the symbol

used. Comparisons are made with the results from a full MD simulation of the same case

Fig. 8 Convergence of the multiscale method applied to the straight channel network cases 1 and 2. The horizontal line is the prescribed tolerance, ζ^{tol}

for case 1. There is quick and accurate convergence of our GeN-IMM solution to the full MD result, and good numerical stability. The error bars for both the full and multiscale solutions are 1.96 standard deviations either side of the mean to represent the 95% confidence interval and take into account the amount of correlation that occurs within each micro-element or full network.

Complete results demonstrating convergence speed and precision for both cases are detailed in Fig. 8, and Tables 1 and 2, respectively. In the tables, the reference mass flow rate $\dot{m}_R$ and reference pressure P_R are taken to be the largest mass flow rate and pressure measured in the multiscale simulations, respectively, as it is these values which most strongly govern the flow characteristics. In both cases the final error in the multiscale solution is $\approx 1\%$ at each boundary for both mass flux and pressure. This can be

Table 1 Straight channel iteration 5: steady-state mass flow rate data for the multiscale method and the full MD solution. Relative error calculated as $|\dot{m}_F - \dot{m}_i|/|\dot{m}_R| \times 100$

Boundary #	Full MD $\dot{m}_F$ (ng/s)	Multiscale $\dot{m}_i$ (ng/s)		Relative error e_i (%)	
		Case 1	Case 2	Case 1	Case 2
1	1.5389	1.5227	1.5457	1.04	0.45
2	1.5387	1.5226	1.5456	1.04	0.45
3	1.5389	1.5310	1.5517	0.51	0.83
4	1.5389	1.5310	1.5517	0.51	0.83
5	1.5389	1.5461	1.5342	0.46	0.30
6	1.5387	1.5461	1.5342	0.48	0.29

Table 2 Straight channel iteration 5: steady-state pressure data for the multiscale method and the full MD solution. Relative error calculated as $|P_F - P_i|/|P_R| \times 100$

Boundary #	Full MD P_F (MPa)	Multiscale P_i (MPa)		Relative error e_i (%)	
		Case 1	Case 2	Case 1	Case 2
1	648.01	647.99	647.60	0.00	0.06
2	606.07	604.81	604.53	0.20	0.25
3	606.07	605.70	604.90	0.06	0.19
4	318.77	318.63	315.96	0.04	0.88
5	318.77	317.68	316.04	0.33	0.86
6	295.06	295.28	293.70	0.07	0.46

compared to relative errors of 73 and 50 % for the initial mass flow rate/pressure predictions in case 1 and 2, respectively. The fluctuation of the mass flow rate between iterations 2 and 5 in Fig. 7a is of similar magnitude to the error bars and so can be accounted for by the noise inherent in the MD simulations.

Setting a fairly tight convergence tolerance, $\zeta^{\text{tol}} = 0.02$, both cases converge comfortably within three iterations. Case 1 may converge slightly quicker than case 2; this is because mass flow rate measurements are more subject to noise than pressure measurements, so we may have a little less confidence in the accuracy of the mass flow rate constraint.

To ensure the multiscale solution is accurately representing the full MD system, longitudinal and transverse profiles are also plotted. Figure 9 shows the measured streamwise pressure and density profiles for case 1 at iteration 5. It is clear that the profiles for the junction micro-elements at the inlet and outlet (micro-elements #1 and #3, respectively) are in good agreement with the results from the full MD simulation. As outlined in Sect. 2.1, in the channel micro-element (#2), a body force is applied over the MD domain that is proportional to the required pressure gradient, and the boundary pressures are then extrapolated from the central point in the channel. As such, in this multiscale solution, the pressure profile in the channel section is linear. Evidently, from Fig. 9, this is not so for the full MD simulation—although it seems to have had little impact on the overall accuracy of the method, suggesting that the linear approximation is acceptable, at least in this case.

Figure 10 presents streamwise velocity and density profiles transverse to the flow for case 1 at iteration 5, superimposed onto the profiles for the full MD simulation. These profiles are measured at the half way point along the network, at the plane A–A (see Fig. 10a). Once again, excellent agreement between the simulations is found, including the accurate capturing of non-equilibrium and

non-continuum effects, namely velocity slip and molecular layering in the velocity and density profiles, respectively.

While accuracy is important, hybrid methods also need to show improved computational efficiency over a full simulation. We measure computational speed-up by comparing the product of the total number of time-steps and the average simulation time for one time-step, for both the full and the hybrid solutions, τ_F and τ_M , respectively. The total number of time-steps was considered to be the number of time-steps taken for the hybrid solution to exhibit the same level of error as the full network solution, and the time taken to reach steady state is included within the measurement time. For a fair comparison, the average simulation time for one time-step was calculated by running each micro element or full network on one central processing unit (CPU). The full MD simulation was run for approximately 1,400,000 time-steps, costing 9.4 s in computational time per time-step. The micro elements for the hybrid solution were run for between 800,000 and 1,200,000 time-steps per iteration, costing between 0.2 and 0.65 s per time-step. Assuming three iterations to convergence, the speed-up τ_F/τ_M is calculated to be 3.9. This is smaller than the value of 7.6 calculated by Borg et al. (2013a) for a similar network because this latter paper combined the two reservoir micro-elements into one, to avoid having to use artificial regions. While this is a useful technique, it is highly specific to the geometry of the network, requiring a serial network with junction components at both the inlet and the outlet. Here, we have relaxed this requirement in order to construct a more general technique. In Borg et al. (2013a), the channel

Fig. 9 Full MD and case 1 multiscale longitudinal profiles of (a) pressure and (b) density (averaged over the cross section). The multiscale density profile in the channel section is generated with the equation of state described in Sect. 2.3

Fig. 10 **a** Schematic showing location of measurement of fluid transverse profiles, at plane A–A. Full MD and case 1 multiscale transverse profiles of **(b)** streamwise velocity, and **(c)** density

height is also slightly greater, so more is gained through the length-scale separation in component #2 than in the present study.

Although the speed-up is fairly modest it should be noted that this test case is developed in order to determine the accuracy of the method. A very large network would have rendered a full MD simulation computationally intractable. So, the computational speed-up is expected to increase dramatically as the scale separation and size of the network increases.

3.2 A bifurcating channel geometry

The second network we test is a bifurcating channel, which demonstrates the generality of the method: this is a configuration that has not previously been open to multiscale solution. The network, along with its hybrid decomposition and full MD setup (for validation purposes), is shown in Fig. 11. The network is split into four components: three channel components at the inlets/outlets and one bifurcating junction component linking them. The lengths of the entrance/exit channel sections of the Y-junction micro-element are ≈ 6 nm in both the real and artificial regions. This size has been chosen conservatively to be at least four times the largest channel flow entrance length calculated from laminar macroscopic flow theory for these cases (≈ 0.35 – 1.5 nm), and to be greater than the MD development length calculated by Borg et al. (2013a), for similar channel cross sections, using a root mean

square deviation approach (≈ 2 – 4 nm). This ensures that the flow in the connecting parts of the Y-junction is fully developed and so does not introduce any artificial disturbances into the multiscale model. Each channel component has a length of $L_1 = L_2 = L_3 = 68$ nm, while their MD micro-elements in the hybrid method are once again $L'_1 = L'_2 = L'_3 = 4.08$ nm, producing length ratios of $L/L' = 16.7$. To add to the complexity and make the solution less apparent, the channel micro-elements are of different heights: 4.08, 2.72, and 3.40 nm for micro-elements #1, #3, and #4, respectively. All micro-elements and the full MD network have a depth in the z -direction of 5.44 nm. This is smaller than in the straight channel network of Sect. 3.1 so that the full MD simulation for validation is not too time-consuming, as the network is now much larger.

Using this network, we modelled two further cases: case 3, which has an inlet at boundary #1 565 MPa, and outlets at boundaries #7 135 MPa and #9 140 MPa; and case 4, which has inlets at boundaries #1 427 MPa and #7 699 MPa, and an outlet at boundary #9 46 MPa. In this way, we simulate two different case variants: a bifurcating case (case 3), with one branch splitting into two; and a mixing case (case 4), with two branches joining into one, see Fig. 11a. For both cases, pressure boundaries are employed at all inlets/outlets.

For both cases, convergence occurred within 3 iterations to a tolerance $\zeta^{\text{tol}} = 0.02$. The mass flow rate and pressure data for case 3 are shown in Tables 3 and 4, respectively. Again, good accuracy is seen, with the mass flow rate

Fig. 11 **a** Schematic of the bifurcating channel multiscale network and its two flow cases with **(b)** the hybrid GeN-IMM decomposition and **(c)** the full MD setup. Dimensions are in nanometres. The boundary numbers are also labelled in each image

Table 3 Case 3 (bifurcating configuration) iteration 5: steady-state mass flow rate data for the multiscale method and the full MD solution. Relative error calculated as $|\dot{m}_F - \dot{m}_i|/|\dot{m}_R| \times 100$

Boundary #	Full MD $\dot{m}_F$ (ng/s)	Multiscale $\dot{m}_i$ (ng/s)	Relative error e_i (%)
1	2.3942	2.3627	1.32
2	2.3942	2.3627	1.32
3	2.3942	2.3535	1.70
4	0.9258	0.9115	1.55
5	1.4691	1.4416	1.88
6	0.9258	0.9133	1.36
7	0.9258	0.9133	1.35
8	1.4691	1.4465	1.54
9	1.4691	1.4464	1.54

errors being less than 2 %, and the pressure errors being mostly less than 1 %, compared to an initial prediction error of 31 %. In addition, the nanoscale fluid phenomena (i.e. velocity slip and molecule layering) seen in the full MD simulation are accurately replicated in all the multiscale model channel micro-elements and at the internal boundaries of the Y-junction micro-element, as shown in Fig. 12.

Tables 5 and 6 show the data for case 4. In this case, the final errors are around 4 % for the mass flow rate and 1.5 % for pressure, compared to an initial prediction error of 53 %. The likely reasons for this larger error can be seen from the longitudinal pressure and density profiles in Fig. 13. While in MD simulations a high pressure gradient is needed to overcome thermal fluctuations, large pressure

Table 4 Case 3 (bifurcating configuration) iteration 5: steady-state pressure data for the multiscale method and the full MD solution. Relative error calculated as $|P_F - P_i|/|P_R| \times 100$

Boundary #	Full MD P_F (MPa)	Multiscale P_i (MPa)	Relative error e_i (%)
1	564.84	564.46	0.07
2	332.28	334.01	0.31
3	332.28	333.86	0.28
4	262.71	267.86	0.91
5	257.16	262.01	0.86
6	262.71	268.57	1.04
7	135.38	135.75	0.07
8	257.16	261.89	0.84
9	139.64	139.53	0.02

Table 5 Case 4 (mixing configuration) iteration 5: steady-state mass flow rate data for the multiscale method and the full MD solution. Relative error calculated as $|\dot{m}_F - \dot{m}_i|/|\dot{m}_R| \times 100$

Boundary #	Full MD $\dot{m}_F$ (ng/s)	Multiscale $\dot{m}_i$ (ng/s)	Relative error e_i (%)
1	1.3179	1.4165	3.85
2	1.3179	1.4165	3.85
3	1.3179	1.4166	3.86
4	-1.2381	-1.2437	0.22
5	2.5565	2.6606	4.07
6	-1.2381	-1.2412	0.12
7	-1.2381	-1.2412	0.12
8	2.5565	2.6486	3.60
9	2.5565	2.6486	3.60

Table 6 Case 4 (mixing configuration) iteration 5: steady-state pressure data for the multiscale method and the full MD solution. Relative error calculated as $|P_F - P_i|/|P_R| \times 100$

Boundary #	Full MD P_F (MPa)	Multiscale P_i (MPa)	Relative error e_i (%)
1	427.02	427.40	0.05
2	313.43	306.56	0.98
3	313.43	306.02	1.06
4	351.41	341.09	1.49
5	218.86	211.13	1.10
6	351.41	340.09	1.62
7	699.15	699.08	0.01
8	218.86	211.22	1.09
9	45.78	45.89	0.02

differences can lead to nonlinearities in pressure and density along the channel. The net flow conductance through the network in case 4 is lower than it is in case 3, so in

order to achieve relatively noise-free mass flow rate measurements, a larger pressure drop is required. This, in turn, creates greater pressure nonlinearity. In MD, pressure nonlinearity arises from varying viscous pressure losses due to the density-dependent viscosity. The rate of viscosity variation increases with increasing density (and thus pressure), and the larger pressure drop in case 4 leads to higher pressures: there are no other pressures as large as those in micro-element #3 in case 4—reaching up to 700 MPa. Furthermore, the more complex the network, the more components are required in a multiscale model; small errors will accumulate as fewer components are directly constrained by the boundary conditions.

The speed-ups for cases 3 and 4 are relatively low. The full MD simulation was run for approximately 1,400,000 time-steps, costing 14.3 and 13.4 s per time-step for cases 3 and 4, respectively. The micro-elements for the hybrid solution were run for 800,000–1,200,000 time-steps per iteration, costing 0.2–3.0 s per time-step. Therefore, $\tau_F/\tau_M = 2.1$ for case 3 and $\tau_F/\tau_M = 2.0$ for case 4. The complex geometry shows the disadvantage of using PBCs in MD. Although physically simple, the necessity for a mirroring boundary leads to a large artificial region that is costly to simulate, especially across multiple iterations. The exploitation of length-scale separation had to be minimised in these test cases in order that we would also be able to simulate the full network (for validation purposes) within a reasonable time.

It should be noted that, although the method is demonstrated here using MD as the micro solver, it is potentially compatible with any other micro solver, such as the direct simulation Monte Carlo (DSMC) method for gas flows, where NPBCs are better developed. Using a different micro solver may therefore remove the need for an artificial region and dramatically increase the speed-up in networks such as this.

4 Conclusion

We have presented a Generalised-Networks IMM (GeN-IMM) that enables the flow of compressible fluids within complex, non-serial nanoscale geometries to be accurately and efficiently modelled. Molecular dynamics (MD) has been used as the micro solver, while the conservation of mass and the continuity of pressure between components provides the macro solution. The advantage of this approach is that nanoscale effects such as slip and molecular layering can be accurately captured within the multiscale solution, whereas a conventional Navier-Stokes-Fourier solution of the network would not be able to predict these important effects. The solver coupling is through the exchange of mass flow rate and pressure information from

Fig. 12 a Schematic showing location of measurements of fluid transverse profiles. Case 3 (bifurcating configuration) full MD and multiscale transverse profiles of streamwise velocity and density at

plane A–A (b, e), plane B–B (c, f), and plane C–C (d, g), and at internal transverse profiles. Case 3 (bifurcating configuration) full MD and multiscale transverse profiles of streamwise velocity and density at

micro to macro, and body force and density controls from macro to micro.

In our network test cases, we found this multiscale method converged after 3 iterations to mass flow rate and pressure errors of $< 4\%$ and usually better than 2% , respectively (when compared to full MD simulations of the same network). The multiscale approach provided a computational speed-up over full MD of between 2 and 4 times.

The new method has some clear advantages over full molecular simulations: (1) it is more efficient than a full MD simulation, and the computational speed-up will be even greater for larger networks (the longer the channels represented by shorter periodic channels, the greater the savings); (2) it can be used as a design tool, due to the ability to set mass-flow-rate conditions as constraints to the solution; (3) it is ideally suited to be run on a small cluster

of CPUs/GPUs (either simultaneously or sequentially, if resources are limited) since the micro-elements are relatively small simulations that can be run independently at each iteration.

The drawbacks are that the method is currently limited to isothermal and steady flows, and that the speed-up in complicated networks is reduced due to the necessary presence of artificial regions in junction micro-elements so that periodic boundary conditions can be employed. However, it should be noted that as the size and complexity of the network increases, the relative size of the artificial regions decreases, and so greater savings are made. We conservatively suggest, as a rule of thumb, that artificial regions should be the same size as the real regions in junction micro-elements, and that each region should have a development length of at least four times that predicted

Fig. 13 Case 4 (mixing configuration): full MD and multiscale longitudinal profiles of (a) pressure, and (b) density. The multiscale density profile in the channel sections are generated with the equation of state described in Sect. 2.3

Table 7 Initial flow conductance coefficients $K_{ij,q}$

Case #	Component # q	Boundary # i, j	Flow coeff. $K_{ij,q}$ ($\times 10^{-20}$ m s)
1 & 2	1	1, 2	−4.2866
	2	1, 2	−1.0708
	3	1, 2	−6.2154
3 & 4	1	1, 2	−1.7281
	2	1, 2	−1.4744
	2	1, 3	−2.0921
	2	2, 3	−0.6242
	3	1, 2	−7.2492
	4	1, 2	−1.1581

by laminar macroscopic flow theory in order to ensure that no disturbance propagates into the flow field and the flow in channel components is fully developed.

Future work should extend the method to quasi-steady problems with large time-scale separation, and to non-isothermal problems. In addition, extending the method to accommodate multispecies flows could lead to interesting mixing cases for which there is currently no multiscale technique. The application of the method in conjunction with other molecular solvers (e.g. DSMC) should also be investigated and, finally, the use of non-periodic boundary conditions could provide additional speed-up for more complex networks as costly artificial regions would not then be required.

Acknowledgments This work is financially supported in the UK by EPSRC Programme Grant EP/I011927/1, and EPSRC Grants EP/K038664/1 and EP/K038621/1. Our full MD calculations were performed on the high performance computer ARCHIE at the University of Strathclyde, funded by EPSRC Grants EP/K000586/1 and EP/K000195/1. **Open Access** This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

Appendix 1: Estimation of flow conductance coefficients $K_{ij,q}$

The flow-conductance coefficients, $K_{ij,q}$ ($=K_{ji,q}$), can be either estimated (e.g. using a continuum model or from past experience of the component geometry type) or calculated using MD simulations from previous iterations and/or those performed prior to the start of the main simulations. The latter is done by solving a set of simultaneous equations for each component q , using pressure and mass flow rate data at boundaries i and j from MD simulations (denoted by $\prime$):

Table 8 Equation of state coefficients $a_{b,q}$

b	$a_{b,q}$
10	-1.0214×10^{-83}
9	4.6745×10^{-74}
8	-9.2942×10^{-65}
7	1.0536×10^{-55}
6	-7.5148×10^{-47}
5	3.5142×10^{-38}
4	-1.0905×10^{-29}
3	2.2310×10^{-21}
2	-2.9569×10^{-13}
1	2.5750×10^{-5}
0	−66.6108

$$\langle \dot{m}_i \rangle' = \sum_{j=1(\neq i)}^W \langle P_i - P_j \rangle' K_{ij,q}. \quad (17)$$

In a component with W boundaries, there are $W(W-1)/2$ independent flow-conductance coefficients $K_{ij,q}$ and each MD simulation provides $W-1$ unique equations. This means $\lceil W/2 \rceil$ MD simulations are required to adequately define all flow conductance coefficients.

In this paper, a mixture of preliminary MD simulations and estimations from experience were used for initial approximations, then data from the latest iteration was used to continuously refine the prediction of $K_{ij,q}$. The initial $K_{ij,q}$ values used are shown in Table 7. In general, the convergence characteristics have been shown to be fairly insensitive to the values of $K_{ij,q}$, but a detailed sensitivity study is necessary future work.

Appendix 2: Equation of state coefficients $a_{b,q}$

The equation of state coefficients $a_{b,q}$ that are used in Eq. (7) are shown in Table 8 for liquid argon, using $\beta = 10$.

Appendix 3: Generating pressures in junction micro-elements

The most convenient and common way of emulating a pressure gradient in simple MD flow channels is to apply a uniform body force to all liquid molecules in the channel (Koplik et al. 1988; Travis et al. 1997; Travis and Gubbins 2000; Fan et al. 2002; Docherty et al. 2014), accompanied by periodicity in the flow direction. However, when the geometry is non-uniform in its cross section, as is the case for

junction micro-elements (and the full MD simulations of the networks), the pressure gradient becomes a varying field. As such, the flow generated by a uniform body forcing will no longer be hydrodynamically equivalent to that created by an imposed pressure difference over the same geometry.

Zhu et al. (2002, 2004) overcome this problem by applying a step body force in only localised regions of uniform cross section. We use the same approach here by introducing an artificial region attached to the real region in the micro-element. Uniform body forces are applied only in the artificial region, in order to impose the correct pressure differences over the real region. This method also retains the simplicity of using periodic boundary conditions (PBCs) in all directions in our simulations. The main drawback of the method is that a complicated component geometry requires a similarly complicated artificial region in the micro-element to ensure the geometrical match necessary for PBCs to function (see Fig. 5). This makes the micro-element larger than it would otherwise be and can have a detrimental effect on overall computational efficiency.

The magnitude of the localised step body force $F_{i,q}$ applied at the i th boundary of the q th component is chosen such that it creates a momentum flux (a pressure jump $\Phi_{i,q}$) within that region that is equal and opposite to that of the pressure difference ΔP we wish to induce over the real region, i.e:

$$-\Delta P = \Phi_{i,q} = F_{i,q} \rho_n \Delta x, \quad (18)$$

where Δx is the extent of the localised region and ρ_n is the number density in this region. Equation (18) can be rearranged in terms of the desired uniform forcing magnitude:

$$F_{i,q} = \frac{\Phi_{i,q}}{\rho_n \Delta x}. \quad (19)$$

References

- Alexiadis A, Lockerby DA, Borg MK, Reese JM (2013) A Laplacian-based algorithm for non-isothermal atomistic-continuum hybrid simulation of micro and nano-flows. *Comput Methods Appl Mech Eng* 264:81–94
- Allen MP, Tildesley DJ (1987) *Computer simulation of liquids*. Clarendon Press, Oxford
- Asproulis N, Kalweit M, Drikakis D (2012) A hybrid molecular continuum method using point wise coupling. *Adv Eng Softw* 46(1):85–92
- Berendsen HJC, Postma JPM, Gunsteren WF, Dinola A, Haak JR (1984) Molecular dynamics with coupling to an external bath. *J Chem Phys* 81(8):3684–3690
- Borg MK, Macpherson GB, Reese JM (2010) Controllers for imposing continuum-to-molecular boundary conditions in arbitrary fluid flow geometries. *Mol Simul* 36(10):745–757
- Borg MK, Lockerby DA, Reese JM (2013a) A hybrid molecular-continuum simulation method for incompressible flows in micro/nanofluidic networks. *Microfluid Nanofluid* 15(4):541–557
- Borg MK, Lockerby DA, Reese JM (2013b) A multiscale method for micro/nano flows of high aspect ratio. *J Comput Phys* 233:400–413
- Borg MK, Lockerby DA, Reese JM (2014) The FADE mass-stat: a technique for inserting or deleting particles in molecular dynamics simulations. *J Chem Phys* 140
- Delgado-Buscalioni R, Coveney PV (2003a) Continuum-particle hybrid coupling for mass, momentum, and energy transfers in unsteady fluid flow. *Phys Rev E* 67(4):1–13
- Delgado-Buscalioni R, Coveney PV (2003b) USHER: an algorithm for particle insertion in dense fluids. *J Chem Phys* 140(2):978–987
- Docherty SY, Nicholls WD, Borg MK, Lockerby DA, Reese JM (2014) Boundary conditions for molecular dynamics simulations of water transport through nanotubes. *Proc Inst Mech Eng Part C J Mech Eng Sci* 228(3):186–195
- E W, Ren W, Vanden-Eijnden E (2009) A general strategy for designing seamless multi-scale methods. *J Comput Phys* 228(15):5437–5453
- Fan X, Phan-Thien N, Yong NT, Diao X (2002) Molecular dynamics simulation of a liquid in a complex nano channel flow. *Phys Fluids* 14(3):1146–1153
- Gad-El-Hak M (2006) *MEMS: introduction and fundamentals*. CRC Taylor and Francis
- Hadjiconstantinou N, Patera A (1997) Heterogeneous atomistic-continuum methods for dense fluid systems. *Int J Mod Phys C* 8(4):967–976
- Huang C, Gregory JW, Sullivan JP (2007) Microchannel pressure measurements using molecular sensors. *J Microelectromech Syst* 16(4):777–785
- Jiang H, Weng X, Li D (2011) Microfluidic whole-blood immunoassays. *Microfluid Nanofluid* 10(5):941–964
- Koplik J, Banavar JR, Willemsen JF (1988) Molecular dynamics of poiseuille flow and moving contact lines. *Phys Rev Lett* 60(13):1282–1285
- Macpherson GB, Reese JM (2008) Molecular dynamics in arbitrary geometries: parallel evaluation of pair forces. *Mol Simul* 34(1):97–115
- Mantzalis D, Asproulis N, Drikakis N (2011) Filtering carbon dioxide through carbon nanotubes. *Chem Phys Lett* 1(3):81–85
- Mattia D, Gogotsi Y (2008) Review: static and dynamic behavior of liquids inside carbon nanotubes. *Microfluid Nanofluid* 5(3):289–305
- Mohamed KM, Mohamad AA (2010) A review of the development of hybrid atomistic-continuum methods for dense fluids. *Microfluid Nanofluid* 8(3):283–302
- Nie XB, Chen SY, E W, Robbins MO (2004) A continuum and molecular dynamics hybrid method for micro- and nano-fluid flow. *J Fluid Mech* 500:55–64
- O'Connell ST, Thompson PA (1995) Molecular dynamics-continuum hybrid computations: a tool for studying complex fluid flows. *Phys Rev E* 52(6):R5792–R5795
- Patronis A, Borg MK, Lockerby DA, Reese JM (2013) Hybrid continuum-molecular modelling of multiscale internal gas flows. *J Comput Phys* 255:558–571
- Praprotnik M, Site LD, Kremer K (2005) Adaptive resolution molecular-dynamics simulation: changing the degrees of freedom on the fly. *J Chem Phys* 123(224):106
- Rapaport DC (2004) *The art of molecular dynamics simulation*. Cambridge University Press, Cambridge
- Ren W, E W (2005) Heterogeneous multi-scale method for the modeling of complex fluids and micro-fluidics. *J Comput Phys* 204(1):1–26
- Thompson PA, Troian SM (1997) A general boundary condition for liquid flow at solid surfaces. *Nature* 389(6649):360–362

- Travis KP, Gubbins KE (2000) Poiseuille flow of lennard-jones fluids in narrow slit pores. *J Chem Phys* 112(4):1984–1994
- Travis KP, Todd BD, Evans DJ (1997) Poiseuille flow of molecular fluids. *Phys A Stat Theor Phys* 240(1–2):315–327
- Wagner G, Flekkøy E, Feder J, Jøssang T (2002) Coupling molecular dynamics and continuum dynamics. *J Comput Phys Commun* 147(1–2):670–673
- Yarin LP, Mosyak A, Hetsroni G (2009) Fluid flow, heat transfer and boiling in micro-channels. Springer, Berlin
- Yasuda S, Yamamoto R (2008) A model for hybrid simulations of molecular dynamics and computational fluid dynamics. *Phys Fluids* 20(11):101–113
- Zhu F, Tajkhorshid E, Schulten K (2002) Pressure-induced water transport in membrane channels studied by molecular dynamics. *Biophys J* 83(1):154–160
- Zhu F, Tajkhorshid E, Schulten K (2004) Theory and simulation of water permeation in aquaporin-1. *Biophys J* 86(1):50–57