THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

The atomistic-continuum hybrid taxonomy and the hybrid-hybrid approach

Citation for published version:

Alexiadis, A, Lockerby, DA, Borg, MK & Reese, JM 2014, 'The atomistic-continuum hybrid taxonomy and the hybrid-hybrid approach', *International Journal for Numerical Methods in Engineering*, vol. 98, no. 7, pp. 534-546. <https://doi.org/10.1002/nme.4646>

Digital Object Identifier (DOI):

[10.1002/nme.4646](https://doi.org/10.1002/nme.4646)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

International Journal for Numerical Methods in Engineering

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

The atomistic-continuum hybrid taxonomy and the hybrid-hybrid approach

Alessio Alexiadis^{a1}, Duncan A. Lockerby^b, Matthew K. Borg^c and Jason M. Reese^d

^a School of Chemical Engineering, University of Birmingham, B15 2TT, United Kingdom

^b School of Engineering, University of Warwick, CV4 7AL, United Kingdom

^c Mechanical & Aerospace Engineering, University of Strathclyde, G1 1XJ, United Kingdom

^d School of Engineering, University of Edinburgh, EH8 9Y, United Kingdom

Abstract

The hybrid taxonomy — a means of characterizing different atomistic-continuum methods on the basis of the type of information exchanged between the atomistic and the continuum solver — is introduced. The formulation of the taxonomy raises a new hybrid possibility, called a ‘hybrid-hybrid’ method. Some examples of hybrid-hybrid simulations for dense fluids are discussed and validated against full-MD results.

Keywords: microfluidics; nanofluidics; atomistic-continuum hybrid methods; molecular dynamics; computational fluid dynamics.

1. Introduction

In recent years there has been a growing interest in atomistic-continuum hybrid modelling of dense fluid systems [1]; this particularly suited for micro- and nano-fluidics problems. For example, the current authors introduced the Laplacian method [2], which is a general framework capable of handling a range of transport phenomena at the micro- and nano-scale.

In this paper we introduce the hybrid taxonomy, which allows us to classify existing hybrid methods based on both the (derivative) order of the information extracted from the microscopic cells and the way this information is handled by the macroscopic solver. This hybrid taxonomy can be used to hypothesize new hybrid models. We focus, in particular, on the *hybrid-hybrid* approach, where two different hybrid models are combined — or *hybridized* — together.

This paper is organized as follows. First the general formulation of the Laplacian method (LM) for dense fluids is reviewed. The mathematical foundation of the method is described in more detail elsewhere [2], but a brief introduction is provided here for completeness. While in previous studies the LM was mainly applied to fully developed channel flows with 1D symmetry, here the method is extended to more general flows. This requires a few additional developments (e.g. a local constraint, and a double grid), which are discussed. The hybrid taxonomy and the hybrid-hybrid concept are then introduced, and some examples of hybrid-hybrid simulations are considered and validated against full-MD results.

2. A brief introduction to Laplacian Method

In order to explain the Laplacian method (LM), we start from the steady state conservation equation for a generic fluid variable ϕ :

¹ email: a.alexiadis@bham.ac.uk

$$\nabla \cdot (\rho \mathbf{u} \phi) + \nabla \cdot \mathbf{J}_\phi - S_\phi = 0, \quad (1)$$

where ρ is the density, $\mathbf{u}$ the streaming velocity, $\mathbf{J}_\phi$ the diffusive flux of ϕ , and S_ϕ the source term. The variable ϕ can be the velocity $\mathbf{u}$ (momentum conservation), the temperature T (heat conservation) or the concentration c of a certain chemical species (mass conservation).

We express the diffusive flux as

$$\mathbf{J}_\phi = -\Gamma \nabla \phi + \Theta(t, \mathbf{x}), \quad (2)$$

where Γ is the generalized diffusion coefficient (e.g. kinematic viscosity μ , thermal conductivity κ , or molecular diffusion D), $\Gamma \nabla \phi$ is the ‘conventional’ part of the flux and Θ the deviation of the real flux from the conventional one. If we introduce eq. (2) in eq. (1), we obtain

$$\Gamma \nabla^2 \phi = \nabla \cdot (\rho \mathbf{u} \phi) + \nabla \cdot \Theta - S_\phi. \quad (3)$$

We now collect all the terms on the right hand side in a new function

$$\Omega(\mathbf{R}) = \frac{1}{\Gamma} [\nabla \cdot (\rho \mathbf{u} \phi) + \nabla \cdot \Theta - S_\phi], \quad (4)$$

and, therefore, we have

$$\nabla^2 \phi = \Omega(\mathbf{R}), \quad (5)$$

where $\mathbf{R}$ is the macroscopic position vector. At this point, we only need a way to estimate $\Omega(\mathbf{R})$ from (local) microscopic models. The simplest way is to approximate Ω with the average microscopic Laplacian

$$\langle \nabla^2 \phi \rangle_{\mathbf{r}=\mathbf{R}} \cong \Omega(\mathbf{R}), \quad (6)$$

Equation (6) is used everywhere in the numerical domain except at boundary nodes, where the macroscopic boundary conditions at the solid bounding walls should be introduced. We extract this information from the microscopic simulations by calculating the value of ϕ , instead of its Laplacian, near the boundaries. Then the equation,

$$\phi(\partial \mathbf{R}) \cong \langle \phi \rangle_{\mathbf{r}=\partial \mathbf{R}}, \quad (7)$$

where $\partial \mathbf{R}$ indicates the boundary nodes, replaces equation (6) for the nodes located at the walls.

We term the procedure from equation (1) to equation (7) the *poissonization* of the conservation equation. Starting from the general conservation equation (1), we arrive at the Poisson equation (7), in which both the non-homogeneous term and the boundary conditions are derived from microscopic (molecular) simulations. From this point of view, poissonization is an instance of the more general concept of *hybridization*. The continuum conservation equation, in fact, can be manipulated in various different ways and we can choose which terms need to be provided by the microscopic solver. We call the final result of this manipulation the *hybridized* conservation equation. Various forms of hybridized equations are presented in Sections 4 and 5.

The goal of the LM hybrid method is to calculate, on a certain number of numerical nodes, the Laplacian (or the value for the boundary nodes) of ϕ and use this result to calculate the macroscopic profile by means of equations (6) and (7). In doing this, we must ensure consistency between the macro and micro domains. The microscopic simulations — in our case, computational Molecular Dynamics (MD) — should therefore be computed under conditions (of temperature, streaming velocity, density etc.) taken from the continuum solution. This can be achieved by using the continuum solution around each node as boundary conditions for the local MD box (see Figure 1).

Each MD box is divided into three regions: the core, the frame, and the buffer. The core is the central region where the information (in this case the Laplacian) from the MD simulation is extracted and transmitted to the macroscopic model. The frame (colored in grey in Figure 1) is the region where the average microscopic velocities, temperatures, densities etc. are constrained in order to match their continuum counterpart. The buffer region surrounds the frame and does not play any role in the exchange of information between the two solvers, but it allows the solution to relax between the external margins of the computational box and the frame.

There are various ways of constraining the velocity, the temperature, and the concentration in the frame regions. Previously [2,3] we have used *velocity-shifting* for the velocity, *velocity-rescaling* for the temperature, and *molecular insertion/deletion* for the concentration. Here we deal with more generic situations, which require a local constraint. The cases discussed in Sections 7 and 8 only concern momentum conservation so only a velocity constraint is necessary. We use the Maxwell Demon technique proposed by Hadjiconstantinou and Patera [4], which is essentially equivalent to the Andersen thermostat [5] with a non-zero average velocity. Some additional remarks on the temperature constraint are given in our Conclusions section.

3. Numerical Algorithm

The goal of the LM is to provide a method that is comparable to MD in terms of accuracy, but computationally cheaper because fewer molecules need to be simulated. There is, however, a price to pay for achieving this goal. Numerical fluctuations, which are typical of MD simulations, depend on the number of molecules in the system. The larger the number of molecules, the lower the numerical noise. Because in hybrid simulations the total number of molecules is reduced, the signal-to-noise ratio necessarily decreases. This is particularly critical in the LM since it is based on evaluating the second derivative of hydrodynamic variables, which tends to amplify the fluctuations in the variables themselves. We note in passing that the problem of low signal-to-noise ratio also often appears in manufactured MEMS, although due to different reasons.

Some noise-reducing methods have been proposed in our previous works [2,3]. A third and more general option is, instead of using the usual definition of the derivative, employing the so-called Lanczos derivative, viz.

$$f'(x) = \lim_{\epsilon \rightarrow 0} \frac{3}{2\epsilon^3} \int_{-\epsilon}^{\epsilon} \xi f(x + \xi) d\xi, \quad (8)$$

which uses a procedure known as Savitzky–Golay filtering in each MD cell [6]. In Figure 2, we illustrate the case of a single MD node with a nondimensional gravity set to $g=0.1$ and fixed nondimensional wall temperature $T_w=1$ (the temperature inside the box, however, rises above this due to viscous dissipation). The total size of the node is $l=5\times5\times5$ molecular units (m. u.), but the Savitzky-Golay filter only acts on the core region of the MD box. The first row in Figure 2 shows the data (streaming velocity and temperature) produced by the MD simulation, the second row a polynomial (second degree for the velocity, and fourth for the temperature) least squares interpolation of the MD data, and the third row the central longitudinal section of the MD box, with the MD data (black diamonds) and the interpolating polynomial (dotted line). The low-noise Laplacian based on Lanczos’ definition is obtained by calculating, at the centre of the cell, the second derivative of the filtering function.

Once the strategy for extracting the Laplacian from each MD box is established, the hybrid algorithm follows the procedure:

Step 1, an initial continuum solution at the N discretization nodes of the macroscopic domain is assumed.

Step 2, knowing the macroscopic solution at the nodes, we can determine N sets of boundary conditions to be applied in the constraint region (the frame) of the microscopic cells.

Step 3, N microscopic simulations are run.

Step 4, the molecular results are processed (Savitzky-Golay filter) in order to extract the local Laplacian and the boundary conditions.

Step 5, equation (6) is solved with the proper boundary conditions (equation 7) and a new macroscopic profile is calculated.

Step 6, repeat from Step 2 until convergence.

The final issue that needs to be discussed concerns Step 2 and, in particular, how to calculate the new boundary conditions for, in our case, MD (i.e. constrained values of ϕ in the frame region) from the macroscopic profile produced in Step 5 (or in Step 1). Here, we introduce a new procedure called the *double macroscopic grid*. Instead of solving equation (6) using the values of the Laplacian $\nabla^2\phi$ extracted at each MD node, we determine a polynomial least-square approximation of these Laplacians based on the Savitzky-Golay filter. Once we have a continuum function approximating $\nabla^2\phi$ over the whole of the macroscopic domain, we are not limited to solving equation (6) only at the N locations of the MD nodes, but we can use a finer grid when calculating ϕ . The advantage of this strategy appears when, from the macroscopic ϕ , we derive the MD boundary conditions (see Figure 3). As an example, suppose that we are solving together the momentum and heat transport equations. In that case the variables under consideration are the streaming velocity u and the temperature T . Beginning at Step 5 (or Step 1) of the algorithm above, we have already obtained the macroscopic T and u profiles on the fine grid. Then we superpose the fine grid on the coarser one. The MD cells, and therefore all the constraint regions (the frames), lie on the coarse grid. By projecting the position of each molecule in the frame to the fine grid, we can determine the streaming velocity and the temperature at the location

occupied by the molecule. With these values, we can then use the Maxwell Demon process in order to assign a velocity to the molecule. Test cases showing how this algorithm works in practice are discussed in Section 6. In the next Section, we discuss a new concept (hybrid taxonomy), which provides a more extensive scope to the procedure.

4. The hybrid taxonomy

We introduce in this section a conceptual classification of different atomistic-continuum hybrid methodologies based on the type of information extracted from the microscopic calculation(s) and on the way this information is managed by the *hybridized* macroscopic equation.

The taxonomy discussed here specifically targets the so-called *embedding-based techniques* [2,7,8] and it does not necessarily apply to other hybrid frameworks, such as the *equation free* [9] or the *domain decomposition method* (DDM) [10]. It is, moreover, only based on the quality of the *microscopic-to-macroscopic* information, while the *macroscopic-to-microscopic* information (i.e. molecular constraint) is not taken into consideration. Our taxonomy, therefore, uses a classification method different from the ‘flux/value’, ‘value/value’, or ‘flux/flux’ typical of DDMs.

Suppose we calculate the macroscopic profile of a certain variable ϕ (e.g. streaming velocity, temperature or concentration) by means of full-MD, full-CFD, or different hybrid methods such as heterogeneous multiscale (HMM) [7], pointwise coupling (PWC) [8], or LM [2] and we compare the results. See Figure 4. Molecular dynamics is a mesh-free method and, therefore, the full-MD results are not concentrated on specific nodes but are available over the whole domain. The other methods, on the contrary, determine the profile by approximating the values of ϕ at certain specific nodes/locations in the domain. In the case of hybrid methods, moreover, the nodes represent microscopic cells, where MD simulations may be run. The information extracted from these MD cells is often in the form of a spatial derivative of ϕ . The HMM or the PWC methods, for instance, use a first order derivative, while the LM uses the Laplacian (a second order derivative). We can also conceive of a 0th-order model. The microscopic solver, for instance, could be used to calculate the value of Γ locally (indicated as $\langle \Gamma \rangle_{\text{MD}}$ in Figure 4 in order to highlight its molecular origin) and subsequently communicate this value to the continuum solver as a position-dependent field.

The role of the macroscopic solver is then to reconstruct/approximate the overall flow profile using the information about the derivative of ϕ at specific locations. If information about all the derivatives is available, we would have a perfect approximation of the full-MD profile. But since we usually only know n derivatives, we only have an n^{th} -order approximation of the profile.

If we do not consider, for the moment, the issue of numerical noise, a higher order approximation should generally indicate higher accuracy. For instance, using HMM to solve the Poiseuille flow problem would require at least ten MD elements to accurately resolve the variation in the gradient of the velocity profile across the channel. However, the LM extracts the 2nd derivative and only one MD element (besides the two at the boundaries) is required to capture the velocity profile. In this sense, the LM has higher accuracy owing to its higher derivative order.

Although this procedure is reminiscent of a Taylor expansion, this is not how the approximation actually works. We do not use a Taylor series to reconstruct the macro field, but a particular hybridized form of the conservation equation (e.g. equation 6 in the case of the LM). ‘Order of approximation’, therefore, does not refer to a formal Taylor expansion. We do not, in fact, necessarily need all derivatives up to the n^{th} to have an n^{th} order approximation. The LM, for instance, tell us that if we poissonize the conservation equation we only need the 2nd derivative, while the 1st derivative is not necessary. We can, nevertheless, design other 2nd order methods, where the continuum equation is manipulated/hybridized in a different way and both 1st and 2nd derivatives are required. We can also have 3rd, 4th or higher order methods. We can, for instance, take the curl of the momentum conservation equation (making it 3rd order, as is done to generate the vorticity transport equation) or re-express it in stream-function form (which is 4th order). As already mentioned, high order derivatives do tend to amplify the numerical noise. For this reason, we have an upper limit for n in practice. Considering the current state-of-the-art for dealing with the issue of numerical noise in atomistic-continuum hybrid methods, this limit seems to be $n=2$ currently.

The prototype equations on the right of Figure 4 are all possible hybridizations of the conservation equation. They are expressed in this particular form to highlight the connection between the macroscopic equation and the (derivative) order of the variable calculated by the microscopic solver (e.g. MD), but there are other options of course. In the case of the LM (2nd order), for instance, instead of the equation indicated in Figure 4, equations (6) and (7) are used. (The equation next to the full-MD picture in Figure 4 indicates that the microscopic solver is used for the entire domain.)

In the hybrid taxonomy, in summary, what we extract from the microscopic domains tells us the order of approximation (not in a formal sense, but as described above), while the final hybridized form of the conservation equation tell us how the n derivatives are combined together. Higher orders have accuracy and efficiency advantages, but carry a noise penalty.

5. The hybrid-hybrid approach

With the hybrid taxonomy discussed in the previous section, we can now hypothesize new multi-scale approaches. So far, we have only considered the conservation equation of a generic variable ϕ , which can be the velocity (momentum transfer), the temperature (heat transfer) or the concentration (mass transfer). We do not need to use the same hybridization, however, for all the fluid variables. A fluid flow system, for instance, may require the Laplacian method for the calculation of the velocity profiles, but a 0th order approximation may be good enough for its temperature profiles.

The idea of what we term the *hybrid-hybrid* approach is to use different hybrid methods for different fluid variables. In this way, a hybrid model is hybridized with another model. The hybridizing model can be another hybrid model with a different order, or even a full-CFD model. According to the nomenclature illustrated in Figure 5, the term *hybrid-hybrid* should be used only for a model combining a hybrid model with full-CFD. A model made of two hybrid models, strictly speaking, should be called a *hybrid-hybrid-hybrid*. In order to avoid redundancy, however, we indicate all these models simply as hybrid-hybrid.

In the case of the calculations considered in the following section, for instance, we use a 2nd order hybrid model for the velocity but a full-CFD solution for the temperature and the pressure/density. The temperature, therefore, is calculated only using the continuum solver and then introduced as a (constant) target temperature for the thermostat applied in the MD simulations. The temperature is not constant in the macroscopic domain, but it is in each MD cell. The pressure should be calculated from the macroscopic continuity equation (mass conservation) — most often through the state equation and/or a procedure like the SIMPLE algorithm [11]. The pressure gradient can be introduced into the MD simulation as a pseudo gravity-like force.

In the examples considered in Section 7 and Section 8 below, the pressure is constant and this last step is not necessary. Some heat, however, is generated by viscous dissipation and the local density slightly changes with temperature. As a consequence, the density in each molecular box is determined beforehand by the equation of state for Lennard-Jones liquids [12].

So far, we have not considered time-scales. We have implicitly assumed that the microscopic processes are considerably faster than the macroscopic one. The macro/micro time-scales, therefore, are completely separated by a sort of molecular quasi-steady assumption. This is the reason why, in Figure 4, the time derivative term is never extracted from the microscopic solver. In practice, all the models represented in Figure 4 can be seen as hybrid-hybrid models in the sense that they are full-CFD in time and n^{th} order in space. This is usually a good approximation in a large variety of practical cases. Hypothetically, however, the time derivative of ϕ could be extracted from the MD calculations too. This means that we can, in principle, couple an n^{th} order scheme in space with a 1st order hybrid scheme in time.

6. Case study

As a test case, we consider an incompressible fluid in a square micro-channel with side L , set in motion by a gravity-like body force g acting in the x -direction (Figure 6). The flow is at steady-state and fully developed. Initially, we run two cases: $g = 0.005$ and $g = 0.01$ (all variables are expressed in standard molecular reduced units). The fluid in the channel is a Lennard Jones (LJ) liquid with density $\rho=0.8$ at $T=1$. Since we use reduced units, the LJ intermolecular potential parameters of the fluid molecules are $\sigma=1$ and $\varepsilon=1$. The side of the square section is $L=106$, the coarse grid is composed of 25 cubic MD simulator cells, 5 in the y -direction and 5 in the z -direction ($N=5 \times 5$). All the MD cells are the same size, $l=5 \times 5 \times 5$ molecular units (m.u.). The thickness of both the frame and the buffer regions in these cells is $l/13$.

We use a hybrid-hybrid algorithm: a 2nd order LM is used for momentum transfer, and full-CFD for heat transfer. While the velocity is calculated with the LM, the temperature is calculated with the continuum solver according to

$$\nabla^2 T = -\frac{\mu}{\kappa}(\tau : \nabla \mathbf{u}), \quad (9)$$

where $\tau : \nabla \mathbf{u}$ is the dissipation function, which in this case simplifies to

$$\tau : \nabla \mathbf{u} = \left(\frac{\partial u}{\partial y} \right)^2 + \left(\frac{\partial u}{\partial z} \right)^2. \quad (10)$$

The temperature at the walls of the channel is constant $T_w=1$, but it slightly rises towards the centre due to viscous dissipation. Since the temperature profile is calculated only with the continuum solver, the numerical value of both the viscosity and the thermal conductivity must be provided as input in equation (9). Here we chose $\mu=1.5$ and $\kappa=6.0$, which are reasonable approximations for Lennard Jones fluids without large velocity and temperature gradients [13]. Pressure is assumed constant and the local density determined by the LJ equation of state [12]. Our hybrid-hybrid results are validated against a full-MD simulation, where the whole channel is simulated with MD and stochastic walls [14] are applied at the boundaries.

Since the stochastic walls generate a temperature jump [2], knowing T_w is not sufficient for the full-CFD part of the hybrid-hybrid procedure. While this value represents the temperature of the wall, equation (9) requires as a boundary condition the temperature of the liquid close to the wall T_w^* . Here we choose $T_w^*=0.75$ for $g=0.005$ and $T_w^*=0.8$ for $g=0.01$ (more details are given in Section 7 below).

Figure 7 summarizes the features of the hybrid-hybrid model used in our example. The model is hybrid-hybrid because it combines LM (which is also a hybrid model) for the conservation of momentum and CFD for the conservation of heat. Figure 7 illustrates also the conceptual differences between the two parts of the model. For the momentum equation, we do not need to know in advance either the shear stress or the slip velocity because the microscopic solver generates this information by itself. For the heat conservation, however, we need to provide the continuum solver (i.e. equation 9) with some additional input, such as the values of the viscosity, the thermal conductivity, and the temperature jump at the boundaries. The continuum solver, moreover, assumes a predetermined relation between J_ϕ and $\nabla\phi$ (i.e. the Fourier law of thermal conduction), which the LM however derives from MD data.

The hybrid-hybrid algorithm, in the specific case under study, follows the procedure:

Step 1, an initial continuum solution at the N discretization nodes of the macroscopic domain is assumed *only* for the velocity profile (here, a parabolic profile with $\mu=1.4$). The temperatures are calculated by means of equation (9) with T_w^* (the boundary condition).

Step 2, knowing the overall density and the local temperature, the local density is calculated using the equation of state.

Step 3, at this point, both the temperature and the density at the N nodes of the domain are known. The MD simulations are run at the nodes; the number of molecules at each node is determined by the local density, while the thermostat is set to the local temperature.

Step 4, the molecular results are processed, the local Laplacian and the boundary conditions for the velocity are calculated, and a new velocity profile is generated.

Step 5, the temperatures are calculated again by using equation (9) with the new velocity profile calculated at Step 5.

Step 6, repeat from Step 2 until convergence.

In theory, this procedure should be iterated until a certain convergence criterion is met. Here, in order to monitor the evolution of the solution, we run our simulation for 20 loops, even if the solution does not change further after 12-15 loops, typically. This is a considerable improvement in speed of convergence over previous versions of the LM [2,3] and is mainly due to the introduction of the double grid.

7. Results

In this section, we compare results obtained with hybrid-hybrid calculations to results from full-MD. This is the standard validation method for hybrid models. As this limits us to domains small enough to be calculated with full-MD, unfortunately it does not fully show the computational savings that can be achieved by using a hybrid method in larger domains. The examples considered here are, on average, two to three times faster than the full-MD results but, as shown elsewhere [1], under the right conditions this performance can be improved to be up to 600 times faster.

Figure 8 shows a comparison between MD and hybrid results. Figures in the first row are velocity results when $g=0.005$, the second row is for $g=0.01$, while the figures in the third row show comparisons on a central longitudinal section (on left when $g=0.005$, on the right when $g=0.01$). In the third row, the solid line represents the full-MD results and the black squares the streaming velocity on the fine grid. Figure 8 also indicates the position of the nodes on the coarse grid (small molecular boxes) with respect to the fine grid (the two grids, at these locations, share the same nodes). These boxes represent the MD cells. In this figure they are illustrated to scale, in order to provide a visual estimation of their size with respect to the size of the whole macroscopic domain. The total number of molecules considering the 25 MD nodes of the hybrid solutions is 3125 (this is a conservative estimate because the number of molecules can change slightly due to the non-constant density), while the full-MD solution is for 50,000 molecules. The 16 nodes at the boundaries, moreover, converge to a steady value much faster; after the first two or three hybrid steps, they do not need to be calculated any further. In this way, the number of molecules involved in the hybrid solution reduces to 1125 (again, a conservative estimate) and, consequently, the computational savings further increase.

When using a hybrid-hybrid model, we focus on a specific variable (here, the velocity) and we look for highly accurate results of this variable. We assume that we are investigating a system where the other variables (here the temperature) are not central and a lower accuracy is acceptable. Despite this, Figure 9 shows that the comparison between the full-MD results and equation (9) are reasonable. This is likely because we chose a LJ fluid in which the temperature gradients are not very high and Fourier's law represents a reasonable approximation. We also had good estimates for μ , κ (from [13]) and T_w^* (from [2]), but this is not always the case.

8. Number of nodes versus size of the nodes

If we double the size of the channel to $L = 212$, and reduce the gravity-like force to $g = 0.0025$ (in order to avoid excessive viscous dissipation), the previous 5×5 grid (labelled mesh A in Figure 10) does not converge to the MD results, as Figure 10 (white circles) indicates.

In order to improve the accuracy, we can increase the number of grid nodes; this is analogous to the typical continuum numerical solver practice of improving the results by using a finer

(macroscopic) grid. Initially, we double the number of nodes of the coarse grid in each direction (i.e. to $N = 10 \times 10$)², but not the size of the MD cells ($l = 5 \times 5 \times 5$ m.u.). This improves the results somewhat (labelled mesh B in Figure 10), but there is still a certain discrepancy between the MD and the hybrid calculations. We could further refine the grid but, at this point, the computational time required for the hybrid calculation matches that of the full-MD and, therefore, it makes little sense to continue in this direction.

In hybrid calculations, however, there is another option to improve the results: we can increase the size, instead of the number, of the molecular cells. Therefore we return to the $N = 5 \times 5$ continuum grid, but now the size of each node is increased to $l = 5 \times 10 \times 10$ m.u. (labelled mesh C). The size of the MD boxes is doubled only in the directions that actually cover the domain. Along the x -coordinate, where the flow is periodic, the size can remain the same. Figure 10 (white squares) shows that, in this case, there is a significant improvement, and the hybrid results are very close to the full-MD solution.

The total number of molecules in the full-MD simulation is 200,000, while both mesh B and mesh C involve, overall, 12,500 molecules. We may wonder, why, if the number of molecules is the same, mesh C is not only more accurate but also faster than mesh B. The reason is that larger MD boxes have smaller fluctuations: the hybrid algorithm spends less time in trying to ‘interpret’ the correct value of the Laplacian at each node and is, therefore, more efficient. This fact is consistent with Figure 4, which shows that the higher the hybrid order, the closer the method is to MD. Our case study here shows, in fact, that in the case of the LM, higher accuracy is achieved not so much through using finer meshes (as in continuum modelling), but rather through using larger microscopic simulation boxes (as in molecular dynamics).

9. Conclusions

We have introduced a hybrid taxonomy, and illustrated how hybrid models of different orders can be combined (or hybridized) together. This has enabled us to propose the new concept of *hybrid-hybrid modelling*. The examples we have discussed in Sections 7 and 8 are cases of hybrid-hybrid modelling, where flow velocities are calculated using the Laplacian method (LM), while temperatures and fluid densities are calculated using CFD.

We have also extended the LM proposed in Alexiadis et al. (2013) [2] to more general configurations. In this case, local constraint should be enforced and we used the so-called Maxwell Demon [4]. This approach has worked well in the examples of Sections 7 and 8, because the LM was applied only to momentum conservation. If it is applied to heat transfer, however, the Maxwell Demon can show some issues when used for constraining the temperature in the small frame region of the molecular boxes (see Figure 1). In order to speed-up the hybrid calculations, in fact, we would like to have high collision frequencies between the constraint region and the thermostat. This, however, perturbs the dynamics of the system and sometimes creates an artificial ‘temperature jump’ between the frame and the core of the cell. We have not shown these results here (and this will be the subject of future studies), but using the Lowe-Andersen thermostat [15] may mitigate this difficulty.

² When we indicate the mesh as $N = 5 \times 5$ or $N = 10 \times 10$, we refer to the coarse grid (see Figure 3). When we illustrate the results in Figure 10, however, the markers indicate the values on the fine grid.

Overall, the most important factor concerning the efficiency of the LM is how quickly the algorithm finds a good estimate for the Laplacian. There are many parameters that affect this (e.g. the size of the MD-box, the degree of the approximating polynomials, the frequency at which the MD profiles are recalculated, etc.) and the difference in terms of computational speed between a good and a bad choice of the parameters can be significant. The values of these parameters, in the examples in Sections 7 and 8, have been determined on a trial-and-error basis, but a significant methodological breakthrough in this area might be achievable through Bayesian-based procedures.

10. References

- [1] Mohamed KM, Mohamad AA. A review of the development of hybrid atomistic-continuum methods for dense fluids, *Microfluidics and Nanofluidics* 2010; **8**: 283–302.
- [2] Alexiadis A, Lockerby DA, Borg MK, Reese JM. A Laplacian-based algorithm for non-isothermal atomistic-continuum hybrid simulation of micro and nano-flows, *Computer Methods in Applied Mechanics and Engineering* 2013; **264**: 81–94.
- [3] Alexiadis A, Lockerby DA, Borg MK, Reese JM. An atomistic-continuum hybrid approach for modelling transport phenomena at the micro- and nano-scale, *Proceedings of the ASME 2013 11th International Conference on Nanochannels, Microchannels, and Minichannels* June 16-10, 2013, Sapporo, Japan.
- [4] Hadjiconstantinou NG, Patera AT. Heterogeneous atomistic-continuum representations for dense fluid systems, *International Journal of Modern Physics C* 1997; **8**: 967–976.
- [5] Andersen HC. Molecular dynamics simulations at constant pressure and/or temperature, *Journal of Chemical Physics*, 1980; **72**: 2384–2393.
- [6] McDevitt TJ. Discrete Lanczos derivatives of noisy data, *International Journal of Computer Mathematics* 2012; **89**: 916–931.
- [7] Ren W, E W. Heterogeneous multiscale method for the modelling of complex fluids and micro-fluidics, *Journal of Computational Physics*, 204, 1–26 (2005)
- [8] Asproulis N, Kalweit M, Drikakis D. A hybrid molecular continuum method using point wise coupling, *Advances in Engineering Software* 2012; **46**: 85–92.
- [9] Gear CW, Hyman JM, Kevrekidis PG, Runborg O, Theodoropoulos C. Equation-free, coarse-grained multiscale computation: enabling microscopic simulators to perform system-level analysis, *Communications in Mathematical Sciences* 2003; **1**: 715–762.
- [10] O'Connell ST, Thompson PA. Molecular dynamics-continuum hybrid computations: A tool for studying complex fluid flows, *Physical Review E* 1995; **52**: 5792–5795.
- [11] Patankar SV. *Numerical Heat Transfer and Fluid Flow*, Taylor & Francis, 1980.
- [12] Toxvaerd S, Praestgaard E. Equation of state for a Lennard-Jones fluid, *Journal of Chemical Physics* 1970; **53**: 2389–2392.

- [13] Rapaport DC. *The Art of Molecular Dynamics Simulation*, 2nd ed. Cambridge University Press, 2004.
- [14] Trozzi C, Ciccotti G. Stationary non-equilibrium states by molecular dynamics. II. Newton's law, *Physical Review A* 1984; **29**: 916–925.
- [15] Koopman and Lowe, Advantages of a Lowe-Andersen thermostat in molecular dynamics simulations *Journal of Chemical Physics* 2006; **124**: 2041031–2041035.

FIGURES

Figure 1. The atomistic-continuum hybrid approach.

Figure 2. Example of Savitzky-Golay filtering in the MD cell.

Figure 3. Schematics of the double macroscopic grid.

full MD

$$\left\langle \frac{\partial \phi}{\partial t} + \Gamma \nabla^2 \phi + S_\phi \right\rangle_{\text{MD}} = 0$$

⋮

2nd order

$$\frac{\partial \phi}{\partial t} + \langle \Gamma \nabla^2 \phi \rangle_{\text{MD}} + S_\phi = 0$$

1st order

$$\frac{\partial \phi}{\partial t} + \nabla \cdot \langle \Gamma \nabla \phi \rangle_{\text{MD}} + S_\phi = 0$$

0th order

$$\frac{\partial \phi}{\partial t} + \langle \Gamma \rangle_{\text{MD}} \nabla^2 \phi + S_\phi = 0$$

full CFD

$$\frac{\partial \phi}{\partial t} + \Gamma \nabla^2 \phi + S_\phi = 0$$

Figure 4. The hybrid taxonomy according to the hybridization order; the $\langle \cdot \rangle_{\text{MD}}$ terms are extracted from the microscopic solver (e.g. MD).

Figure 5. Illustration of the hybrid-hybrid nomenclature.

Figure 6. The square channel configuration under study.

Figure 7. Schematic representation of the hybrid-hybrid scheme used in the case study.

Figure 8. Full-MD vs. hybrid (fine grid) velocity results.

Figure 9. Full-MD vs. continuum (equation 9) temperature results in a central section of the channel (fine grid).

Figure 10. Comparison among three different hybrid meshes (central section).