

Original citation:

Lockerby, Duncan A., Patronis, Alexander, Borg, Matthew K. and Reese, Jason M.. (2014) Asynchronous coupling of hybrid models for efficient simulation of multiscale systems. *Journal of Computational Physics*, 284 . pp. 261-272.

Permanent WRAP url:

<http://wrap.warwick.ac.uk/77101>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work of researchers of the University of Warwick available open access under the following conditions.

This article is made available under the Creative Commons Attribution 4.0 International license (CC BY 4.0) and may be reused according to the conditions of the license. For more details see: <http://creativecommons.org/licenses/by/4.0/>

A note on versions:

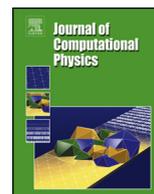
The version presented in WRAP is the published version, or, version of record, and may be cited as it appears here.

For more information, please contact the WRAP Team at: publications@warwick.ac.uk

warwick**publications**wrap

highlight your research

<http://wrap.warwick.ac.uk>



Asynchronous coupling of hybrid models for efficient simulation of multiscale systems



Duncan A. Lockerby^{a,*}, Alexander Patronis^a, Matthew K. Borg^b,
Jason M. Reese^c

^a School of Engineering, University of Warwick, Coventry CV4 7AL, UK

^b Department of Mechanical & Aerospace Engineering, University of Strathclyde, Glasgow G1 1XJ, UK

^c School of Engineering, University of Edinburgh, Edinburgh EH9 3JL, UK

ARTICLE INFO

Article history:

Received 6 August 2014

Received in revised form 16 December 2014

Accepted 20 December 2014

Available online 24 December 2014

Keywords:

Multiscale simulations

Unsteady micro/nano flows

Hybrid methods

Scale separation

Rarefied gas dynamics

ABSTRACT

We present a new coupling approach for the time advancement of multi-physics models of multiscale systems. This extends the method of E et al. (2009) [5] to deal with an arbitrary number of models. Coupling is performed asynchronously, with each model being assigned its own timestep size. This enables accurate long timescale predictions to be made at the computational cost of the short timescale simulation. We propose a method for selecting appropriate timestep sizes based on the degree of scale separation that exists between models. A number of example applications are used for testing and benchmarking, including a comparison with experimental data of a thermally driven rarefied gas flow in a micro capillary. The multiscale simulation results are in very close agreement with the experimental data, but are produced almost 50,000 times faster than from a conventionally-coupled simulation.

© 2014 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

A multi-physics description of a multiscale system is often referred to as a ‘hybrid’ model. In fluid dynamics, a typical hybrid combines a molecular treatment (a ‘micro’ model) with a continuum-fluid one (a ‘macro’ model), with the aim of obtaining the accuracy of the former with the efficiency of the latter [1–4]. The micro and macro models generally have characteristic timescales that are very different, which means that time-accurate simulations can be extremely challenging: the size of the timestep required to make the micro model stable and accurate is so small that simulations over significant macro-scale time periods are intractable. If the system is ‘scale-separated’, a physical (as distinct from numerical) approximation can be made that enables the coupled models to advance at different rates (asynchronously) with negligible penalty on macro-scale accuracy. E et al. [5] were the first to introduce and implement this concept in a time-stepping method for coupled systems, referred to in the classification of Lockerby et al. [6] as a continuous asynchronous (CA) scheme (‘continuous’ since the micro and macro models advance without interruption [5]). In this paper we extend this idea to multiscale systems comprising an arbitrary number of coupled models.

* Corresponding author.

E-mail address: duncan.lockerby@warwick.ac.uk (D.A. Lockerby).

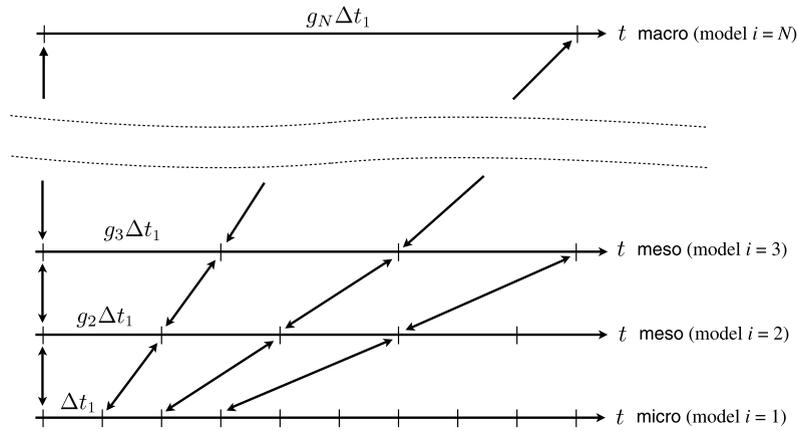


Fig. 1. The continuous asynchronous (CA) coupling scheme extended to multi-model multiscale systems.

2. Extension to multi-model systems

We consider an N -model timescale-separated system, where the i th model has a characteristic timescale T_i , and indexing is ordered such that

$$T_i \leq T_{i+1} \quad \text{for } i = 1, \dots, N - 1. \tag{1}$$

Model $i = 1$ is the micro model and $i = N$ is the macro model; models $i = 2$ to $i = N - 1$ are ‘meso’ models. The degree of scale separation S_i between model i and $i + 1$ is

$$S_i = \frac{T_{i+1}}{T_i} \geq S_{\text{tol}}, \tag{2}$$

where the tolerance S_{tol} requires each distinct model of the system to be scale separated from every other to some degree, for example, $S_{\text{tol}} = \mathcal{O}(10)$. If this condition is not met, the two models are treated as one, and coupling is performed conventionally.

In general each model can be considered to have its own time variable (t_i), and be represented by

$$\frac{dX_i}{dt_i} = \mathcal{F}_i(X(t_i)), \tag{3}$$

where X_i are the set of variables of the i th model, and \mathcal{F}_i is some function of the complete system’s variables, $X = \{X_1, X_2, \dots, X_N\}$. It is important to make clear the distinction between the characteristic timescale of the i th model in isolation (T_i) and the timescale of its variables within the coupled system; they are, potentially, completely different.

To solve this set of models, the independent time variables must be related to each other. If all time variables are equal (i.e. $t = t_{1\dots N}$) the system is conventionally coupled. However, we can advance models at different rates with the physical modification

$$t_1 = t_2/g_2 = \dots = t_N/g_N, \tag{4}$$

where g_i is the rate that the i th model advances relative to the micro model. This approximation provides a means to exchange fine timescale resolution for long timescale predictions, and the extent to which it is valid depends on the degree of scale separation between models, i.e. on the magnitude of S_i . For coupled models that are highly scale separated ($S_i > S_{\text{tol}}$), the smaller-scale model will remain quasi-equilibrated to the dynamics of the larger-scale model despite the physical modification, and so behave similarly to as in the unmodified system. The aim is thus to represent the scale-separated system ($> S_{\text{tol}}$) with one that is less, but still significantly, scale separated ($= S_{\text{tol}}$): this is how acceptable values of g_i are determined (see below for the specific procedure). Detailed analyses of the error associated with this physical approximation for a two-model system are given in E et al. [5] and Lockerby et al. [6].

Fig. 1 provides an illustration of a numerical implementation of Eq. (4) using different timestep sizes for each model, while exchanging variables as if the timesteps were equivalent (this is *asynchronous* coupling). The timestep of the i th model is

$$\Delta t_i = g_i \Delta t_1, \tag{5}$$

where Δt_1 is the micro model timestep.

The aim of this asynchronous coupling scheme is to maximise the total simulated period. This is done by maximising the timestep in each model subject to the following constraints:

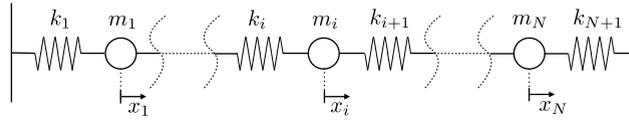


Fig. 2. A coupled mass–spring system.

1. The physical approximation of Eq. (4) represents a very scale-separated system by one that is less, but still to a degree, scale separated (i.e. has a scale separation of S_{tol}). This places an upper limit on the amount a timestep of one model can be increased relative to another:

$$\Delta t_i \leq \frac{S_{i-1}}{S_{\text{tol}}} \Delta t_{i-1}.$$

2. The numerical accuracy is satisfactory and stability guaranteed for each individual model:

$$\Delta t_i \leq \Delta t_{i,\text{max}},$$

where $\Delta t_{i,\text{max}}$ is an estimation of the maximum timestep that is permissible for each model.

Based on these constraints we can set the timestep of each model recursively:

$$\Delta t_i = \min \left\{ \Delta t_{i,\text{max}}; \frac{S_{i-1}}{S_{\text{tol}}} \Delta t_{i-1} \right\}, \quad \text{for } i = 2, \dots, N, \tag{6}$$

where $\Delta t_1 = \Delta t_{1,\text{max}}$.

We now consider a series of examples. The examples of Sections 3–5 are used to illustrate the effectiveness, challenges, and shortcomings of employing the physical approximation in Eq. (4) and applying time stepping of Eq. (6); the example of Section 6 provides a demonstration of an application to hybrid continuum-molecular modelling. Note, there are a range of numerical methods that could be applied to the numerically-stiff systems of Sections 3–5, but which cannot be applied to the hybrid example of Section 6.

3. Example 1: A simple mass–spring system

A serial mass–spring system with N masses and $N + 1$ springs is shown in Fig. 2. The governing equations for the i th model are

$$\frac{d}{dt} \begin{bmatrix} v_i \\ x_i \end{bmatrix} = \begin{bmatrix} \frac{1}{m_i} (k_i (x_{i-1} - x_i) + k_{i+1} (x_{i+1} - x_i)) \\ v_i \end{bmatrix}, \tag{7}$$

where x_i and v_i are the displacement and velocity of the i th mass (m_i) and k_i is the i th spring constant. A characteristic timescale for each model can be obtained from its natural period in isolation:

$$T_i = 2\pi \sqrt{\frac{m_i}{k_i + k_{i+1}}}. \tag{8}$$

In this example $N = 5$, the spring constants are equal, and the $(i + 1)$ th mass is $900\times$ heavier than the i th mass, such that

$$T_{i+1} = 30T_i. \tag{9}$$

All models are thus significantly scale-separated.¹

The complete system of equations is solved using the midpoint method (a second-order Runge–Kutta scheme), but with timesteps for each model chosen according to Eq. (6), with $\Delta t_{i,\text{max}} = T_i/100$.

In the first case, the initial displacements $x_{i,t=0}$ are chosen such that, when the masses are released from rest, only the slowest eigenmode of the system is excited. Fig. 3(a) shows the response of the lightest (m_1) and heaviest (m_5) masses governed by the micro and macro models, respectively; the analytical eigenvalue solution (the dashed line) is indistinguishable from the multiscale solution. Fig. 3(b) shows the solution using $S_{\text{tol}} = 5$, which places a less conservative restriction on the balance between accuracy and efficiency. For $S_{\text{tol}}=10$, the model advances $81\times$ faster than if standard (synchronous) coupling were used; for $S_{\text{tol}} = 5$ it advances $1296\times$ faster.

The masses are now initially displaced an equal distance and then released from rest; this excites all eigenmodes, to a degree. Fig. 4 shows: (a) the response of a meso model ($i = 3$) and (b) the response of the macro model ($i = 5$, the heaviest mass), for $S_{\text{tol}} = 5$. The macro description is accurate, while for the meso model only the low frequency eigenmode is accurately captured. This example illustrates the fundamental trade-off required in multiscale modeling: efficiency in predicting macro variations can be dramatically increased, but at the expense of micro/meso scale resolution.

¹ For clarity, we have chosen characteristic timescales of the models to increase with i so that Eq. (1) is satisfied without having to change index notation.

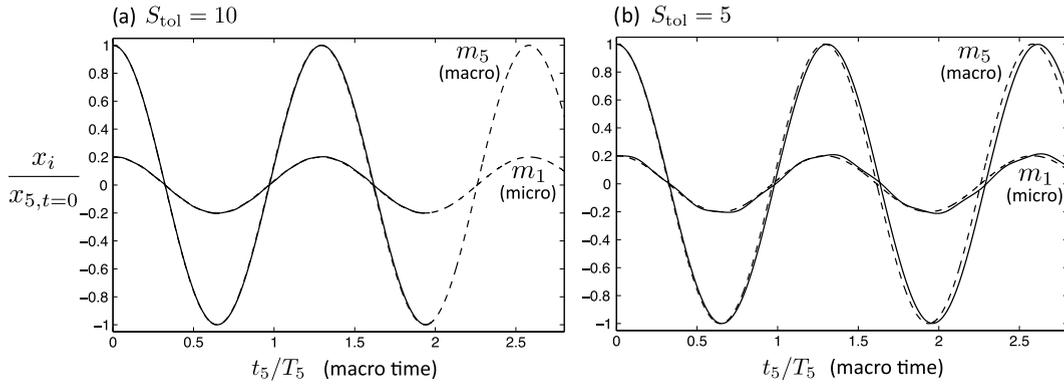


Fig. 3. Normalised displacement response of the lightest and heaviest masses (m_1 and m_5 , respectively) to excitation of the slowest eigenmode. The multiscale result (–) and an analytical eigenvalue solution (––).

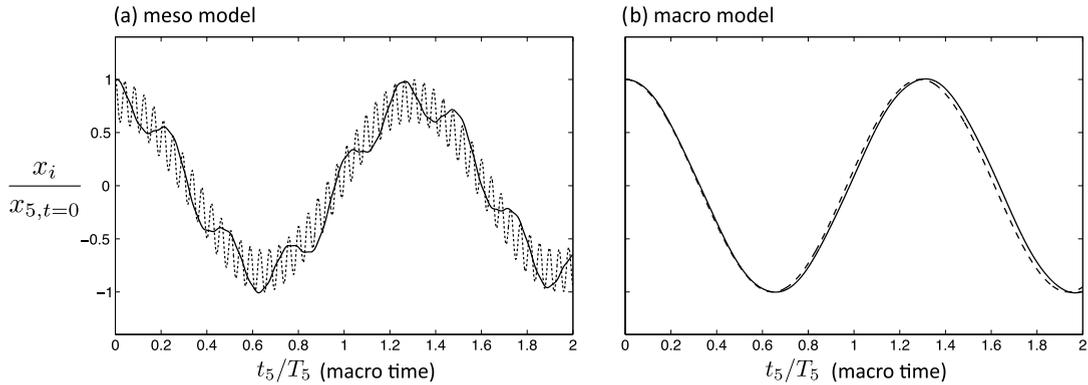


Fig. 4. Normalised displacement response of (a) the median and (b) the heaviest masses (m_3 and m_5 , respectively) to an equal initial displacement. The multiscale result (–) and an analytical eigenvalue solution (––).

4. Example 2: A Lotka–Volterra system

The Lotka–Volterra equations, in various forms, have been applied to an extremely diverse range of problems, spanning economics [7], biology [8] and chemistry [9]. Originating from the analysis of auto-catalytic chemical reactions [9], the equations are now most commonly used to study the population dynamics of competing biological systems, which is the example we consider here.

The population growth rate of a species in a (sequential) food chain is as follows:

$$\frac{dy_i}{dt} = y_i(-r_i + p_i y_{i+1} - q_i y_{i-1}), \tag{10}$$

where y_i is the population size of the i th species, r_i is the intrinsic death rate (in the absence of any prey or predator), p_i is the population growth rate due to the consumption of lower species in the food chain, and q_i is the death rate due to predation from higher species in the food chain. The intrinsic death rate of each species defines a characteristic timescale for that species model ($r_i = 1/T_i$), and which classifies the model in the macro-to-micro hierarchy; this rate increases moving up the food chain. Thus, the Apex Predator, at the top of the food chain, has the highest intrinsic death rate (in the absence of prey), and its population size, y_1 , is governed by the micro model.

Table 1 gives parameters for a food-chain example consisting of four species. For the initial population sizes Y_i the ecosystem is in equilibrium, and the numbers of each species will remain constant. If, however, all plants are removed ($Y_4 = 0$), the populations of the remaining species will eventually reduce to extinction. Fig. 5 shows the population response of each species in the ecosystem, as predicted by a standard numerical solution (the dashed line, used as a benchmark) and the multiscale approach using $S_{tol} = 10$ (the solid line). As in Section 3, time integration is performed using a second-order Runge–Kutta method, with timesteps for each model chosen according to Eq. (6), with $\Delta t_{i,max} = T_i/5$. For the benchmark solution $\Delta t_i = \Delta t_{1,max}$.

The first observation is that the slowly varying population sizes of the Apex Predators and the Herbivores are predicted very accurately by the multiscale scheme, which is 100× computationally faster than the benchmark solution. Fig. 6 shows the macro model prediction using less conservative tolerances on the minimum acceptable scale separation, i.e. $S_{tol} = 2.5$ and $S_{tol} = 5$ (which are 1600× and 400× faster to compute than the benchmark, respectively). As expected, the multiscale

Table 1
Lotka–Volterra parameters for a 4-species food chain.

Trophic level	Model	i	Y_i (initial population)	r_i	p_i	q_i
Apex predators	micro	1	10^2	10^4	10^1	0
Predators	meso	2	10^3	10^2	11×10^{-2}	10^1
Herbivores	macro	3	10^4	10^0	2×10^{-5}	10^{-3}
Plants	–	4	10^5	0	0	0

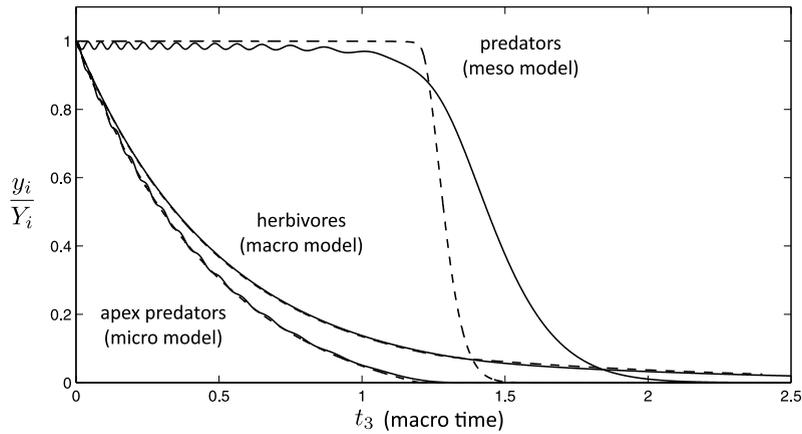


Fig. 5. Normalised population response to the instantaneous removal of Plants. The multiscale result (–) and a conventional (benchmark) numerical solution (––).

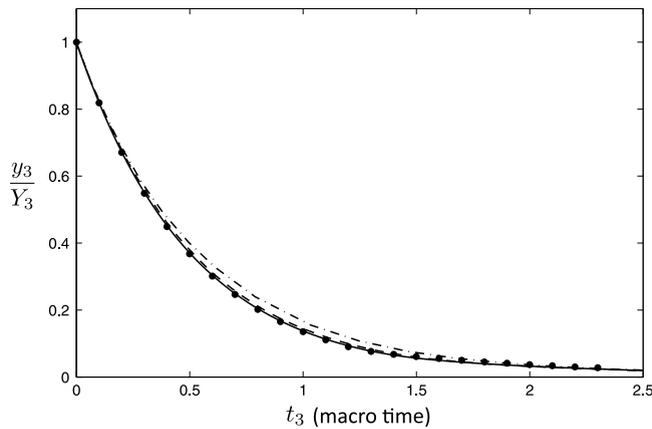


Fig. 6. Normalised Herbivore population response to the instantaneous removal of Plants: $S_{tol} = 10$ (–), $S_{tol} = 5$ (––), $S_{tol} = 2.5$ (–·–), and a benchmark numerical solution (•). Note, for clarity, the benchmark solution is not plotted at every timestep.

solution converges to the numerical benchmark as S_{tol} is increased; $S_{tol} = 10$ appears to provide a very accurate result for the macro variable.

However, compared to the other species, the Predators’ population decline is rapid, and occurs after some delay. The delay occurs because, initially, the Predators’ prey and the Predators’ predators are both reducing – only when their prey is significantly diminished is there a major reduction in Predator numbers. The multiscale approach does not capture, with any fidelity, these shorter scale phenomena, and actually introduces erroneous short timescale oscillations. This again highlights that exploiting scale separation affords very efficient prediction on large timescales, but at the expense of finer timescale resolution. Note, in this case, the micro model prediction is very good, because the short scale response only manifests itself in the meso model’s variable.

5. Example 3: A lubrication system

In this section we consider air-layer lubrication of a liquid journal bearing, as depicted in Fig. 7. The air layer, despite being thin, can significantly reduce the overall drag on the bearing, due to the lower viscosity of air relative to liquids. This simple air-layer lubrication concept is exploited in super-hydrophobic coatings, which have a chemical hydrophobicity and

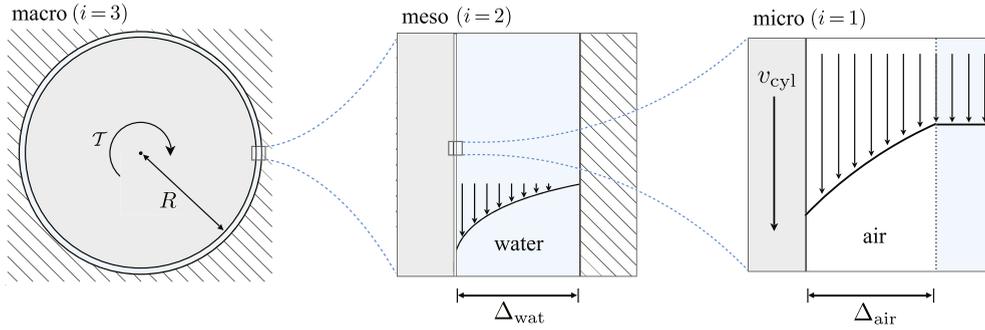


Fig. 7. Schematic of a liquid journal bearing with a lubricating air layer.

surface topology that, when submerged in water, combine to trap air pockets on the surface. Such coatings have applications in marine drag reduction [10,11] and for self-cleaning surfaces [12]. In the context of multiscale modelling they are relevant because of the very different scales associated with the air layer, external water, and the body/vehicle.

The bearing example of this section consists of a steel cylinder, of radius $R = 10$ cm, to which is applied an oscillatory torque, \mathcal{T} . The cylinder rotates within a fixed outer cylinder containing water; the surface of the inner cylinder is coated with an air layer of thickness $\Delta_{\text{air}} = 1$ μm ; and the annular thickness of water is $\Delta_{\text{wat}} = 0.1$ mm (i.e. $R \gg \Delta_{\text{wat}} \gg \Delta_{\text{air}}$); see Fig. 7.

The low-speed, unsteady, incompressible Navier–Stokes equations provide models for the air layer and the water (i.e. the micro and meso models, $i = 1$ and $i = 2$, respectively):

$$\frac{\partial v_{\text{air}}}{\partial t_1} = \frac{\mu_{\text{air}}}{\rho_{\text{air}}} \frac{\partial^2 v_{\text{air}}}{\partial r^2} \quad (\text{micro}) \quad (11)$$

and

$$\frac{\partial v_{\text{wat}}}{\partial t_2} = \frac{\mu_{\text{wat}}}{\rho_{\text{wat}}} \frac{\partial^2 v_{\text{wat}}}{\partial r^2} \quad (\text{meso}) \quad (12)$$

where r is the radial coordinate from the cylinder centre, v is tangential velocity, μ is dynamic viscosity, ρ is density, and the subscripts ‘air’ and ‘wat’ denote the respective fluids. Given that $R \gg \Delta_{\text{wat}} \gg \Delta_{\text{air}}$, the curvature of the geometry can be neglected. The micro and meso models are coupled by the requirement for the shear stress and the velocity to be continuous at the air–water interface (assuming no slip):

$$\mu_{\text{air}} \left. \frac{dv_{\text{air}}}{dr} \right|_{r=r_{\text{int}}} = \mu_{\text{wat}} \left. \frac{dv_{\text{wat}}}{dr} \right|_{r=r_{\text{int}}}, \quad (13)$$

and

$$v_{\text{air}}|_{r=r_{\text{int}}} = v_{\text{wat}}|_{r=r_{\text{int}}}, \quad (14)$$

where the radial position of the air–water interface is $r_{\text{int}} = R + \Delta_{\text{air}}$. The water at the wall of the outer cylinder is stationary (i.e. there is no slip).

Newton’s second law determines the evolution of the tangential velocity of the inner cylinder surface (v_{cyl}); this is the macro model ($i = 3$):

$$\frac{\partial v_{\text{cyl}}}{\partial t_3} = \frac{R}{I} \left[2\pi LR^2 \left(\mu_{\text{air}} \left. \frac{\partial v_{\text{air}}}{\partial r} \right|_{r=R} \right) + \mathcal{T}(t_3) \right] \quad (\text{macro}) \quad (15)$$

where L is the length of the bearing (into the page) and I is the moment of inertia of the cylinder. The applied torque is $\mathcal{T} = A \sin(\omega t_3)$, where ω is the angular frequency and A is the amplitude. The macro model is coupled to the micro model through shear stress in the air layer at the cylinder wall (i.e. the term in parenthesis in Eq. (15)) and via no-slip at the cylinder–air interface:

$$v_{\text{cyl}} = v_{\text{air}}|_{r=R}. \quad (16)$$

See Appendix A for values of the physical parameters used in this example.

The characteristic timescales estimated for each model are the viscous timescale (for the two fluids), and a fraction of the torque period (for the cylinder):

$$T_1 = \frac{\rho_{\text{air}} \Delta_{\text{air}}^2}{\mu_{\text{air}}}; \quad T_2 = \frac{\rho_{\text{wat}} \Delta_{\text{wat}}^2}{\mu_{\text{wat}}}; \quad T_3 = \frac{\pi}{4\omega}. \quad (17)$$

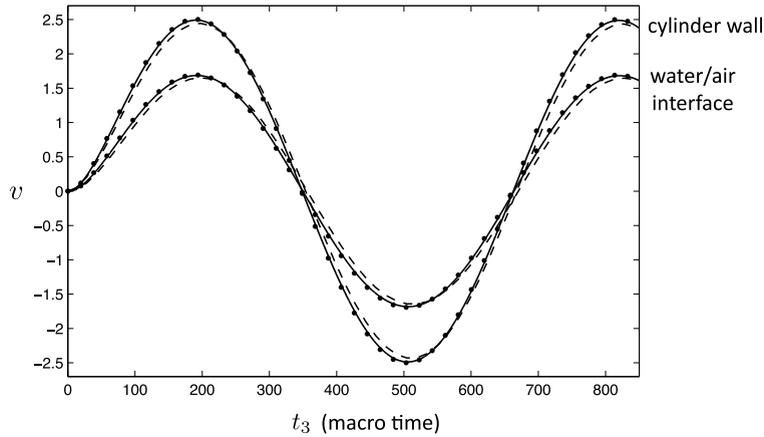


Fig. 8. Tangential velocity v [m s^{-1}] developing in macro time t_3 [s] for the cylinder wall and the air–water interface. Response to an oscillatory cylinder torque. $S_{\text{tol}} = 20$ (\bullet), $S_{\text{tol}} = 10$ ($-$), and $S_{\text{tol}} = 5$ ($-.$). Note, for clarity, the $S_{\text{tol}} = 20$ result is not plotted at every timestep.

In this example, the disparity in timescales is vast: $T_3/T_1 \sim 10^9$. For consistency with previous examples, the midpoint method is used for time-advancement, with timestep sizes determined by Eq. (6) (see Appendix A for $\Delta t_{\text{max},i}$). Spatial discretisation of the fluid models, i.e. of Eqs. (11)–(12), is performed using a second-order central-difference approximation; for this illustrative example, only 10 grid points are used in each fluid layer (a finer mesh does not substantially change the results).

Fig. 8 shows the variation of the velocity of the cylinder wall, and the velocity of the air–water interface, with macro time. The velocity of the cylinder wall is almost 50% higher than that of the air–liquid interface (which would be the approximate velocity of the cylinder wall if no air layer were present); the drag coefficient of the cylinder in water has been reduced by almost 50% due to the presence of the thin air layer.

Here it is not practical to benchmark the multiscale results against a conventional numerical solution (i.e. one with equal timestep sizes), because of the high computational cost to obtain the latter. Instead, and what must be done in practice, is to show the independence of the multiscale result to increases in S_{tol} . This is akin to a grid-dependency study – setting a larger S_{tol} has the effect of reducing the difference between timestep sizes.

Fig. 8 shows results for $S_{\text{tol}} = 5, 10,$ and 20 ; the results for $S_{\text{tol}} = 10$ and 20 are barely distinguishable, indicating that $S_{\text{tol}} = 5$ is a fair prediction, and $S_{\text{tol}} = 10$ is a very accurate one. The computational speed-up afforded by the asynchronous timestep coupling is in this case extremely high: $\times 9.5 \cdot 10^7$ (for $S_{\text{tol}} = 5$); $\times 1.2 \cdot 10^7$ (for $S_{\text{tol}} = 10$); and $\times 3 \cdot 10^6$ (for $S_{\text{tol}} = 20$).

Now we consider the sudden application of a constant torque, $\mathcal{T} = 1$, to the stationary system. This case highlights the potential difficulties in identifying characteristic timescales. The macro model, Eq. (15), does not have an inherent timescale in the absence of an oscillatory torque. In other words, in isolation (i.e. without air or water), the cylinder would perpetually accelerate in response to the constant torque. In these circumstances some estimate of the timescale of the model when interacting with others, is needed. Here we achieve this with an approximation of the acceleration and velocity of the cylinder wall in terms of the air-layer shear stress, and combine them to get a timescale.

In the absence of an applied torque, the acceleration of the cylinder walls will be proportional to the shear stress in the air layer at the wall (τ_{wall}), and inversely proportional to the moment of inertia of the cylinder (see Eq. (15)):

$$\frac{\partial v_{\text{cyl}}}{\partial t} \propto \frac{LR^3 \tau_{\text{wall}}}{I}. \quad (18)$$

If we assume a linear velocity profile in the air layer, the velocity of the cylinder wall will be proportional to the shear stress and the air-layer thickness, but inversely proportional to the dynamic viscosity, i.e.

$$v_{\text{cyl}} \propto \frac{\tau_{\text{wall}} \Delta_{\text{air}}}{\mu_{\text{air}}}. \quad (19)$$

Division of Eq. (19) by (18) gives a characteristic timescale that we can use in our simulation:

$$T_3 = \frac{\Delta_{\text{air}} \rho_{\text{cyl}} R}{\mu_{\text{air}}}, \quad (20)$$

where ρ_{cyl} is the density of the steel cylinder. With the exception of this macro timescale T_3 , and $\Delta t_{3,\text{max}} = T_3/200$, all other parameters are the same as above.

Fig. 9 shows the velocity response of the cylinder wall and air–water interface to the suddenly-applied constant torque. Again, calculations are performed using $S_{\text{tol}} = 5, 10,$ and 20 , with $S_{\text{tol}} = 10$ providing a result that appears to be insensitive to further increases of S_{tol} . This solution is achieved $\times 6.5 \cdot 10^6$ faster than a solution using equal timesteps. The characteristic

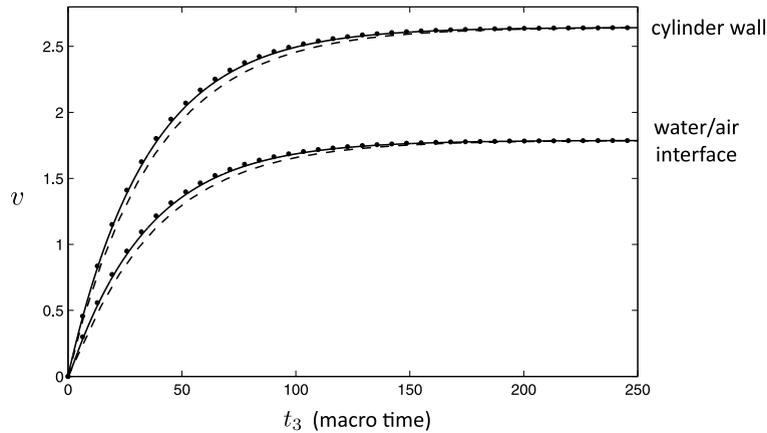


Fig. 9. Tangential velocity v [m s^{-1}] developing macro time t_3 [s] for the cylinder wall and the air–water interface. Response to a suddenly-applied constant cylinder torque. $S_{\text{tol}} = 20$ (\bullet), $S_{\text{tol}} = 10$ (–), and $S_{\text{tol}} = 5$ (–). Note, for clarity, the $S_{\text{tol}} = 20$ result is not plotted at every timestep.

timescale predicted by Eq. (20), $T_3 = 78.5$ s, is reasonable given the observed timescales. Even so, if this prediction had been much different, the main consequence would be that the S_{tol} -independence threshold would be different, and the dependency study might have required additional simulations to find that threshold.

6. Example 4: A Knudsen compressor

Finally, we consider the rarefied gas flow between two reservoirs, held at different temperatures, connected by a thin cylindrical capillary: a single-stage Knudsen compressor, see Fig. 10. Rarefaction effects in the capillary transport gas from the cold to the hot reservoir; this counter-intuitive phenomenon is thermal transpiration (sometimes known as thermal creep) and was first observed by Reynolds [13]. The configuration shown in Fig. 10 was constructed by Rojas-Cárdenas et al. [14] in order to study the transient behaviour of thermal transpiration in a closed system; some of their experimental data is presented below.

In terms of simulation, this system cannot be modelled using standard Navier–Stokes equations and boundary conditions, since thermal transpiration is a *thermodynamic non-equilibrium* phenomenon [15,16]. On the other hand, an accurate gas-kinetic treatment would be computationally intractable over the entire domain. To tackle this, we decompose the system into three coupled models, applying the appropriate modelling assumptions to each: the reservoir model (macro, $i = 3$); the capillary model (meso, $i = 2$); and the gas-kinetic model (micro, $i = 1$); see Fig. 10. The macro model defining the reservoir pressures is obtained from mass conservation and by assuming an ideal gas:

$$\begin{aligned} \frac{dp_c}{dt_3} &= -\frac{\mathcal{R}\theta_c}{V_c} \dot{m}_{(z=0)}, \\ \frac{dp_h}{dt_3} &= -\frac{\theta_h}{\theta_c} \frac{V_c}{V_h} \frac{dp_c}{dt_3}, \end{aligned} \quad (21)$$

where p is pressure, \mathcal{R} is the gas constant, θ is temperature, V is the reservoir volume, \dot{m} is the mass flow rate along the capillary, z is distance along the capillary from the cold to hot reservoir, and the subscripts c and h denote the cold and hot reservoirs, respectively. Note, here we assume that there is no significant change in mass of gas within the capillary, though this can easily be accounted for if necessary.

The meso model for the high-aspect-ratio capillary is obtained from the continuity equation integrated over the cross section:

$$\frac{\partial p}{\partial t_2} + \frac{\mathcal{R}\theta}{A} \frac{\partial \dot{m}}{\partial z} = 0, \quad (22)$$

where A is the cross-sectional area of the capillary. The meso model is coupled to the macro model by the boundary conditions: $p = p_c$, $\dot{m} = \dot{m}_{(z=0)}$ at $z = 0$; and $p = p_h$ at $z = \ell$, where ℓ is the length of the capillary. The temperature variation along the capillary is prescribed (using the same fit as in Rojas-Cárdenas et al. [14]) by

$$\theta = \theta_c + (\theta_h - \theta_c)(e^{\alpha z} - 1)/(e^{\alpha \ell} - 1), \quad (23)$$

where α is a constant.

The micro model, \mathcal{G} , provides a means to close the entire system, by relating mass flow rate to pressure and temperature:

$$\frac{\partial \dot{m}}{\partial t_1} = \mathcal{G}\left(\frac{\partial \theta}{\partial z}; \frac{\partial p}{\partial z}; \mathcal{X}\right), \quad (24)$$

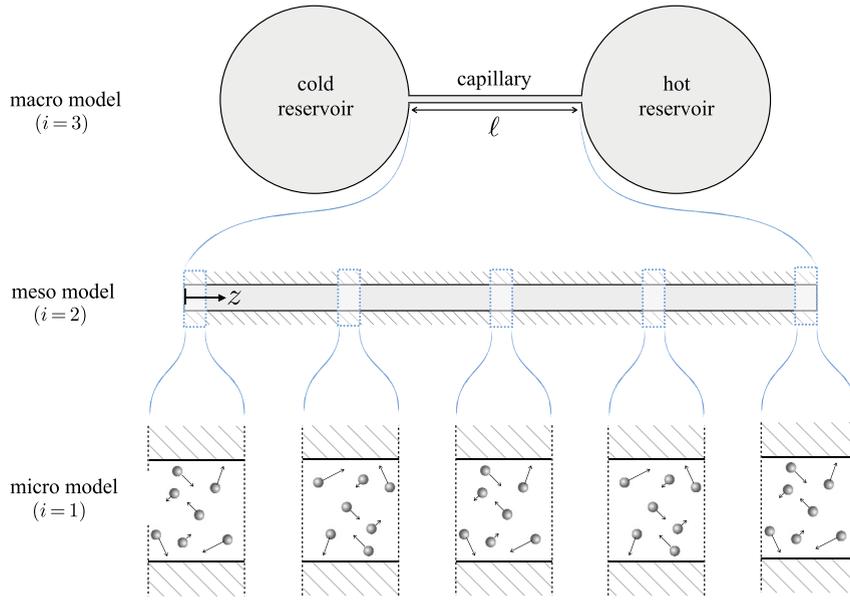


Fig. 10. Schematic of a multiscale simulation strategy for the single-stage Knudsen compressor experimental configuration of Rojas-Cárdenas et al. [14].

where \mathcal{X} contains information regarding the molecular structure of the gas, which is required to accurately model thermal transpiration. Here, the micro model \mathcal{G} is a spatially-distributed array of low-variance deviational simulation Monte Carlo (LVDSMC) subdomains (see Fig. 10). LV-DSMC is a particularly accurate and low noise method for simulating small deviations from equilibrium in rarefied gas flows [17,18]. Using micro particle-simulation subdomains to represent points in the meso domain is substantially more efficient than modelling the entire channel with a single particle simulation – this is an application of the Internal Multiscale Method (IMM), and readers are referred to [19–21] for a detailed description. The simulated particles of each (streamwise periodic) subdomain are forced by an effective body force, which represents the equivalent pressure and temperature gradient occurring at that location in the meso model (the pressure and temperature are also set). The micro model is thus coupled to the meso model by the streamwise pressure gradient, pressure, and mass flow rate at each of the subdomain locations. For the simulations we present here, 12 subdomains are used. For accuracy, the derivatives in z featuring in Eqs. (22) and (24) are evaluated from a Chebyshev polynomial interpolation of p and \dot{m} from subdomain locations corresponding to Chebyshev–Gauss–Lobatto points.

Viscous development within the cross-section of the capillary defines the characteristic timescale of the micro model

$$T_1 = \frac{\bar{\rho} R_{\text{cap}}^2}{\bar{\mu}}, \quad (25)$$

where R_{cap} is the capillary radius and $\bar{\rho}$ and $\bar{\mu}$ are the average initial density and viscosity of the gas. If we assume a quasi-steady velocity profile (which is only valid for $t \gg T_1$), characteristic timescales of the meso and macro model can be estimated from Eqs. (22) and (21), respectively:

$$T_2 = \frac{\bar{\mu} \ell^2}{\bar{p} R_{\text{cap}}^2}, \quad (26)$$

and

$$T_3 = \frac{\bar{\mu} \ell V_t}{\bar{p} R_{\text{cap}}^4}, \quad (27)$$

where \bar{p} is the average initial pressure and V_t is the total volume of the combined reservoirs.

The experiments of Rojas-Cárdenas et al. [14] were performed with Argon gas, a borosilicate (glass) capillary of circular cross-section with $\ell = 52.7 \pm 0.1$ mm, $R_{\text{cap}} = 242.5 \pm 3$ μm connecting two reservoirs of volume $V_c = 19.81 \pm 0.54$ cm^3 and $V_h = 14.85 \pm 0.40$ cm^3 , held at $\theta_c = 301$ K and $\theta_h = 372$ K. A heater applied to the hot reservoir generated a temperature distribution through the capillary fitted by Eq. (23), with $\alpha = 84.82$ m^{-1} . Initially, the two reservoirs were held at fixed pressure ($\bar{p} = 237.7$ Pa), allowing thermal transpiration flow to develop; the system was then closed, and the pressure in the two reservoirs allowed to equilibrate. Note, in a gas that is non-rarefied the initial state would be the equilibrium one.

We ran multiscale simulations of this experimental configuration using Euler time-stepping, for simplicity, with timestep sizes chosen by Eq. (6). Our complete simulation parameters are provided in Appendix B.

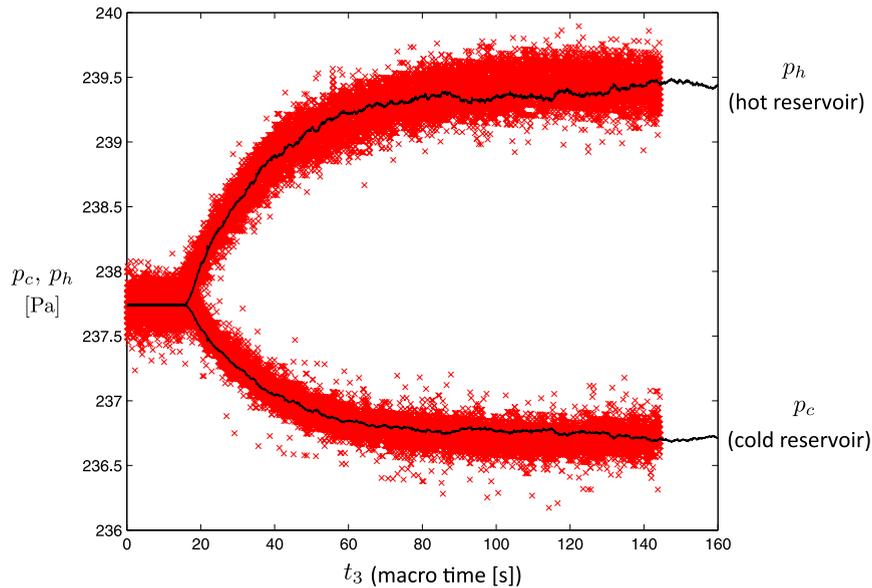


Fig. 11. Comparison of experimental data and the multiscale solution, for transient development of a Knudsen compressor; a plot of reservoir pressures versus time. Both reservoirs are held at a constant pressure until approximately $t_3 = 18$ s, at which point the reservoirs are instantaneously closed to the environment. Experimental data of Rojas-Cárdenas et al. [14] (\times) and the multiscale simulation (—).

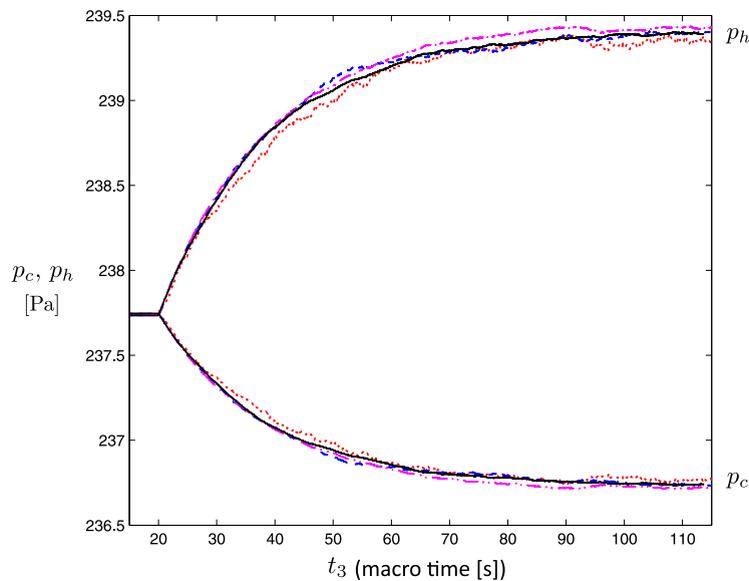


Fig. 12. Reservoir pressure versus time for increasing values of S_{tol} : 10 (\cdots); 20 ($-\cdots$); 30 ($-\cdot-$); 40 (—).

Fig. 11 shows the transient response of the pressure in each reservoir after the system is instantaneously closed; there is very close agreement between the experimental measurements and the multiscale simulation. The asymmetry of the figure around the initial pressure is caused by the different reservoir volumes ($V_c > V_h$). The multiscale result is obtained with $S_{\text{tol}} = 10$, and required the use of twelve Intel Xeon X5650 2.66 GHz cores for 4.13 hours (wall-clock time). If the time steps were set equal (i.e. a conventional synchronous coupling) the simulation would have taken over 20 years on the same hardware. In fact, the saving over conventional modelling is far greater than this, if we also take into account the saving due to *spatial* multiscaling by using the IMM [19–21].

Fig. 12 shows the impact of increasing S_{tol} ; a similar result is obtained in all cases, but with lower noise at higher S_{tol} . This highlights an important general limitation of multiscaling with stochastic models (e.g. LVDSMC) or inherently noisy methods (e.g. Molecular Dynamics): fewer timesteps means less sampling, and thus more noise. The trade-off in hybrid (continuum-particle) multiscaling can thus be summarised:

fine-scale accuracy \leftrightarrow noise & large-scale prediction.

7. Discussion and summary

In this paper we have described how an asynchronous coupling of multiple models can be used to balance fine-scale accuracy with long timescale predictions in scale-separated systems. A physical approximation is made that represents the scale separated system by one that is less, but still to some extent, scale separated. In the examples presented, we have been able to make very large computational savings with very little reduction in macro-scale accuracy.

In the majority of these examples, it has been possible to identify clear characteristic timescales, upon which timestep sizes for each model are then based. In many practical cases, though, such identification will not be straightforward. Multiple time scales will be present, and heavy reliance on crude estimation will be necessary. In these cases, performing a dependency study on S_{tol} is essential for obtaining reliable predictions.

Applying the scheme to sets of models that, despite having disparate characteristic timescales in isolation, are strongly nonlinearly coupled, presents another issue (e.g. chaotic systems). These systems can demonstrate extreme macro-scale sensitivity to small scale fluctuations (i.e. are not separable). Here, again, the minimum test of solution integrity must be that the results are independent beyond a threshold S_{tol} .

Significant future technical development needs are towards coping with systems exhibiting time-varying degrees of scale separation. A fully adaptive and automated scheme – one which does not rely on the prediction of characteristic timescales of models – is the next major step forward.

Acknowledgements

This work is financially supported by EPSRC grants EP/I011927/1, EP/K038664/1, and EP/K038621/1. The authors would like to thank: Marcos Rojas-Cárdenas, Irina Graur, Pierre Perrier, and J. Gilbert Méolans for providing their experimental data; Nicolas Hadjiconstantinou for providing the LVDSMC source code; and Carlos Duque-Daza for the suggestion to consider the Lotka–Volterra equations as an example.

Appendix A. Parameters for simulations of Section 5

$R = 10$ cm; $\Delta_{\text{air}} = 1$ μm ; $\Delta_{\text{wat}} = 0.1$ mm; $\mu_{\text{air}} = 18.6 \times 10^{-6}$ kg m⁻¹ s⁻¹; $\mu_{\text{wat}} = 8.9 \times 10^{-4}$ kg m⁻¹ s⁻¹; $\rho_{\text{air}} = 1.23$ kg m⁻³; $\rho_{\text{wat}} = 1000$ kg m⁻³; $\rho_{\text{cyl}} = 1.23$ kg m⁻³; $L = 1$ m; $I = 1.26$ kg m²; $f = 1.59$ mHz; $A = 1$ Nm; $\Delta t_{1,\text{max}} = T_1/200$; $\Delta t_{2,\text{max}} = T_2/200$; $\Delta t_{3,\text{max}} = T_3/400$.

Appendix B. Parameters for simulations of Section 6

$\ell = 52.7$ mm; $R_{\text{cap}} = 242.5$ μm ; $V_c = 19.81$ cm³; $V_h = 14.85$ cm³; $\theta_c = 301$ K; $\theta_h = 372$ K; $\alpha = 84.82$ m⁻¹; $\bar{p} = 237.7$ Pa; $T_1 = 9.907$ μs ; $T_2 = 4.205$ ms; $T_3 = 47.03$ s; $\Delta t_{1,\text{max}} = 8 \times 10^{-9}$ s (set by best-practice guidelines for LV-DSMC); $\Delta t_{2,\text{max}} = T_2/100$; $\Delta t_{3,\text{max}} = T_3/100$; 12 independent LV-DSMC subdomains are positioned on a Lobatto grid along the capillary; On average $\sim 10,000$ deviational particles are used per subdomain. Purely diffuse reflection is assumed at walls. A Variable Hard Sphere (VHS) model of Argon is used; VHS diameter = 4.17×10^{-10} m; VHS mass = 6.634×10^{-26} kg.

References

- [1] W. E, B. Engquist, X. Li, W. Ren, E. Vanden-Eijnden, Heterogeneous multiscale methods: a review, *Commun. Comput. Phys.* 2 (2007) 367–450.
- [2] M. Kalweit, D. Drikakis, Multiscale simulation strategies and mesoscale modelling of gas and liquid flows, *IMA J. Appl. Math.* 76 (2011) 661–671.
- [3] K.M. Mohamed, A.A. Mohamad, A review of the development of hybrid atomistic-continuum methods for dense fluids, *Microfluid. Nanofluid.* 8 (2010) 283–302.
- [4] M.K. Borg, D.A. Lockerby, J.M. Reese, Fluid simulations with atomistic resolution: a hybrid multiscale method with field-wise coupling, *J. Comput. Phys.* 255 (2013) 149–165.
- [5] W. E, W. Ren, E. Vanden-Eijnden, A general strategy for designing seamless multiscale methods, *J. Comput. Phys.* 228 (2009) 5437–5453.
- [6] D.A. Lockerby, C.A. Duque-Daza, M.K. Borg, J.M. Reese, Time-step coupling for hybrid simulations of multiscale flows, *J. Comput. Phys.* 237 (2013) 344–365.
- [7] R. Goodwin, *A Growth Cycle – Socialism, Capitalism and Economic Growth*, C.H. Feinstein (Ed.), Cambridge University Press, 1967.
- [8] V. Volterra, Variations and fluctuations of the number of individuals in animal species living together, *J. Cons. - Cons. Perm. Int. Explor. Mer* 3 (1926) 3–51.
- [9] A.J. Lotka, Contribution to the theory of periodic reactions, *J. Phys. Chem.* 14 (1910) 271–274.
- [10] R.J. Daniello, N.E. Waterhouse, J.P. Rothstein, Drag reduction in turbulent flows over superhydrophobic surfaces, *Phys. Fluids* 21 (2009) 085103.
- [11] M.A. Samaha, H.V. Tafreshi, M. Gad-el-Hak, Influence of flow on longevity of superhydrophobic coatings, *Langmuir* 28 (2012) 9759–9766.
- [12] B. Bhushan, Y.C. Jung, K. Koch, Self-cleaning efficiency of artificial superhydrophobic surfaces, *Langmuir* 25 (2009) 3240–3248.
- [13] O. Reynolds, On certain dimensional properties of matter in the gaseous state. Parts I and II, *Philos. Trans. R. Soc. Lond.* 170 (1879) 727–845.
- [14] M. Rojas-Cárdenas, I. Graur, P. Perrier, J.G. Méolans, Time-dependent experimental analysis of a thermal transpiration rarefied gas flow, *Phys. Fluids* 25 (2013) 072001.
- [15] M. Gad-el-Hak (Ed.), *MEMS: Introduction and Fundamentals*, CRC Press, 2005.
- [16] J.M. Reese, M.A. Gallis, D.A. Lockerby, New directions in fluid dynamics: non-equilibrium aerodynamic and microsystem flows, *Philos. Trans. R. Soc., Math. Phys. Eng. Sci.* 361 (2003) 2967–2988.
- [17] T.M.M. Homolle, N.G. Hadjiconstantinou, A low-variance deviational simulation Monte Carlo for the Boltzmann equation, *J. Comput. Phys.* 226 (2007) 2341–2358.

- [18] G.A. Radtke, J.-P.M. Péraud, N.G. Hadjiconstantinou, On efficient simulations of multiscale kinetic transport, *Philos. Trans. R. Soc., Math. Phys. Eng. Sci.* 371 (2013) 20120182.
- [19] M.K. Borg, D.A. Lockerby, J.M. Reese, A multiscale method for micro/nano flows of high aspect ratio, *J. Comput. Phys.* 233 (2013) 400–413.
- [20] A. Patronis, D.A. Lockerby, M.K. Borg, J.M. Reese, Hybrid continuum-molecular modelling of multiscale internal gas flows, *J. Comput. Phys.* 255 (2013) 558–571.
- [21] A. Patronis, D.A. Lockerby, Multiscale simulation of non-isothermal microchannel gas flows, *J. Comput. Phys.* 270 (2014) 532–543.