

Article

Transient Two-Way Molecular-Continuum Coupling with OpenFOAM and MaMiCo: A Sensitivity Study

Helene Wittenberg ^{*}  and Philipp Neumann 

High Performance Computing, Helmut-Schmidt-University, Holstenhofweg 85, 22043 Hamburg, Germany; philipp.neumann@hsu-hh.de

^{*} Correspondence: h.wittenberg@hsu-hh.de

Abstract: Molecular-continuum methods, as considered in this work, decompose the computational domain into continuum and molecular dynamics (MD) sub-domains. Compared to plain MD simulations, they greatly reduce computational effort. However, the quality of a fully two-way coupled simulation result strongly depends on a variety of system-specific parameters, and the corresponding sensitivity is only rarely addressed in the literature. Using a state-flux molecular-continuum coupling algorithm, we investigated the influences of various parameters, such as the size of the overlapping region, the coupling time step and the quality of ensemble-based sampling of flow velocities, in a Couette flow scenario. In particular, we considered a big setup in terms of domain size and number of time steps, which allowed us to investigate the long-term behavior of the coupling algorithm close to the incompressible regime. While mostly good agreement was reached on short time scales, it was the long-term behavior which differed even with slightly differently parametrized simulations. We demonstrated our findings by measuring the error in velocity, and we summarize our main observations with a few lessons learned.



Citation: Wittenberg, H.; Neumann, P. Transient Two-Way Molecular-Continuum Coupling with OpenFOAM and MaMiCo: A Sensitivity Study. *Computation* **2021**, *9*, 128. <https://doi.org/10.3390/computation9120128>

Academic Editor: Demos T. Tsahalidis

Received: 27 October 2021

Accepted: 24 November 2021

Published: 30 November 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: coupled simulation; MD; CFD; sensitivity; Couette flow

1. Introduction

Numerical models for fluid flow usually simulate the conditions for one specific scale in time and space. Physical processes, however, may operate on different scales, and numerical models need to be accurate enough to capture the corresponding phenomena. This can be done by explicit modeling of effects that are connected to each other, or by coupling two or even more models of the scales involved (multiscale modeling). An exemplary phenomenon is the flow of a liquid in a nanochannel: while the flow in the bulk of the channel might still be captured with a continuum model, the same might not be possible for the slip of the fluid along the walls, which arises from molecular wall-fluid interactions. One way to solve this in the context of multiscale modeling is to use a molecular-continuum model, which relies on a molecular dynamics (MD) simulation of the fluid close to the wall, a continuum-based computational fluid dynamics (CFD) model in the bulk and the coupling of both models.

A molecular-continuum scheme consists of two simulation models and the coupling framework, which results in a complex system with many physical and technical parameters. Consequently, the number of possible parameter combinations is very high. However, only few studies have reported on parameter sensitivities. This has two effects. First, it makes it impossible to perfectly reproduce results, which means typically being limited to making coarse-scale, statistical quantitative comparisons only, or even making qualitative comparisons only. Second, it is important to know the influences of all parameters and their interconnections to be able to use a molecular-continuum coupled model for the simulation of other arbitrary setups.

The first molecular-continuum coupling approach was developed by O'Connell and Thompson in 1995 [1]. Since then, several research groups published papers discussing

variations, improvements and applications of similar models. Most have concentrated on novel coupling methodologies, and only a few have commented on the impacts of potential coupling parameters [1–12]. A comprehensive review of the general development in the field was given by Mohamed and Mohamad [13]. One of the main distinctions among coupling methodologies is flux-based vs. state-based coupling. For state-based coupling, state variables, such as mass or velocity, are communicated and applied as boundary conditions to the other domain. In contrast, flux-based coupling exchanges the flux variables. This method has been used to capture fluctuations, which can be important on the mesoscale [7], but at the same time it is more expensive for incompressible flows [13] and its stability can be relevant [14]. State-based coupling is more stable [14] and easier to implement [13]. However, calibration of the solvers is often needed to assure agreement in the fluid's transport properties [13]. In the following we present a short summary of research regarding the technical parameters of molecular-continuum coupling.

The sizes of the overlapping region and the non-overlapping molecular-dynamics region were studied by Yen et al. [15]. For their evaluation, a hybrid molecular-continuum model was applied for the simulation of a transient Couette and Poiseuille flow with a slip condition at the wall. A minimum of 12σ pure MD regions was found to be necessary to simulate the correct slip length, where σ is the characteristic length scale at the MD scale, considering a single-site Lennard–Jones model. For liquid argon, σ has a value of roundabout 0.34 nm. For their setup, the authors found a size of the overlap region of at least 10σ to improve the convergence. They pointed out that a certain minimum size of the inner molecular dynamics region is important for valid and correct simulations with the hybrid model.

An even more intensive study on the optimal design of the overlap layer was conducted by Wang et al. [16]. Their research was performed using a hybrid molecular-continuum model, coupling LAMMPS and OpenFOAM, for the simulation of a transient Couette flow with slip and no-slip boundary conditions at the wall. Wang et al. split the overlap layer into different functional regions: control (for boundary forces and mass flow modeling); data transfer from continuum to particle region and vice versa; and a buffer region between the data transfer regions. They evaluated the effect of varying the widths of these regions. The best results were obtained when the continuum-to-particle transfer region was twice as wide as the others. Furthermore, they dealt with the topic of how often data from the macroscopic solver need to be sent to the continuum; and how many sampling points should be included to obtain good agreement between both solvers, but at the same time do as few operations as possible. Their main finding was that instead of using all sampling points for the data exchange, just a few need to be included to obtain almost the same accuracy as when using all. Here, the term sampling points refers to the data of one molecular dynamics time step. For more details, the reader should refer to the original paper.

A molecular-continuum coupled simulation can be performed using various open-source software approaches. Besides general frameworks for the coupling of numerical simulation models—for example, preCICE (<https://github.com/precice/precice>, accessed on 9 November 2021) specific open-source libraries for molecular-continuum coupling have been developed, such as the CPL-library (<https://github.com/Crompulence/cpl-library>, accessed on 9 November 2021), MaMiCo (<https://github.com/HSU-HPC/MaMiCo>, accessed on 9 November 2021) and HAC-par [17]. Since they are customized for these kinds of simulations, they offer many special features (for example, MaMiCo's multi-instance sampling, which means several quasi-identical, yet randomized molecular dynamics simulations are performed, and hydrodynamic quantities are sampled from this ensemble before it is applied in the continuum simulation) and are generally more suitable for molecular-continuum coupling. For example, the mdFoam+ (<https://github.com/MicroNanoFlows/OpenFOAM-2.4.0-MNF>, accessed on 9 November 2021) molecular dynamics model enables the use of OpenFOAM (<https://www.openfoam.com>, accessed on 9 November 2021) for a molecular-continuum coupling [18]. Due to its

integration with the OpenFOAM package, it can be used together with a broad range of solvers, pre-processing tools and post-processing tools. However, due to the file-based data exchange in OpenFOAM, which mdFoam+ inherits, and due to issues with input and output in parallelized simulations, mdFoam+ coupled simulations have drawbacks in terms of their parallelized execution on high performance computing clusters.

We carried out a detailed parameter sensitivity study for transient molecular-continuum coupling. We used the macro-micro-coupling tool (MaMiCo), due to its high number of molecular-continuum specific features, good scalability, modular design and correspondingly high level of configurability. The tool allows one to couple arbitrary MD and CFD solvers. Thus, the implemented coupling algorithm, which is abstracted from the underlying solvers, can be reused by other research groups. Uptake of the methods and their application, however, requires a detailed understanding of the parametrization. Therefore, we demonstrate the impacts of various parameters by comparing the results of coupled molecular-continuum simulations with a reference setup. The evaluation incorporates the quantitative differences in velocity and density fields, and a qualitative analysis of the behavior of the fluid. We aimed to develop an understanding of the influences of the parameters and also explain why a specific parameter leads to a certain result in the simulation. Sub-optimal configurations imply quantitatively small errors. These inconsistencies may not be detectable in small coupled simulations, executed over short time periods. For this reason, we simulated a relatively big scenario consisting of 180,000 particles and executed over 150,000 molecular dynamics time steps. Compared to similar studies in this field [1,15,16,19], our simulation is the biggest one in terms of particle updates, and it is of similar size to others with regard to time steps and spatial dimensions.

Due to the simplicity of its setup, its interesting features (sharp gradients at the beginning, linear flow profile in steady state, etc.) and the existence of an analytical solution, we used start-up Couette flow. For the continuum model, we used a finite-difference solver from the OpenFOAM package. The MD solver SimpleMD from the MaMiCo package was applied, and the coupling of both was accomplished by MaMiCo.

We give an introduction to the coupling method in Section 2.3. The implementation of the coupling of the OpenFOAM continuum solver and MaMiCo is—as this is also a novel development—described in Section 2.4. Afterwards, the simulations and their results are presented in Section 3, followed by a discussion of the lessons we learned during the project regarding coupled molecular-continuum simulations.

2. Materials and Methods

2.1. Molecular Dynamics Simulation

The basic idea of molecular dynamics models [20] is the simulation of individual molecules or particles on the basis of Newton's equation of motion and a model for the inter-molecular forces:

$$\begin{aligned} \frac{d}{dt} \vec{v}_i &= \frac{1}{m} \vec{F}_i \\ \frac{d}{dt} \vec{x}_i &= \vec{v}_i \end{aligned} \quad (1)$$

with \vec{x}_i , \vec{v}_i , \vec{F}_i , m denoting position, velocity, force and mass of a molecule i . We make use of the SimpleMD library [21,22], which is restricted to single-site, truncated shifted Lennard–Jones molecules; note that, due to choosing MaMiCo for coupling, other more flexible MD codes such as LAMMPS could be used as well, and corresponding results of our study immediately carry over to other implementations. For the reduction of the computational costs, the interaction of particle pairs with a distance wider than the cutoff radius is neglected.

2.2. Finite Volume Method and OpenFOAM

The continuum part was modeled using a CFD solver from the library OpenFOAM, an open-source, state-of-the-art CFD framework which offers an extensive range of features and has a large user base from different fields. Due to the broad use, the frequent updates

and the good functionality, it is an interesting enhancement for MaMiCo to offer an interface for OpenFOAM, which was implemented recently and is also detailed in the following.

We chose icoFoam as the solver from OpenFOAM, as it allows, amongst others, for the simulation of geometrically simple, transiently incompressible, laminar flows, such as our Couette flow. For the application of a different OpenFOAM solver coupled with MaMiCo, technical extensions to the interface would be necessary—for example, the consideration of density variations in the case of a compressible solver.

OpenFOAM is written in object-oriented style using C++. The functions and classes of OpenFOAM are designed to be easily extensible and reusable. According to this concept, the solvers consist of stand-alone executables, which call all the necessary OpenFOAM functionalities. To run a simulation using the icoFoam solver, one only needs to execute the icoFoam script. The scrip starts by reading the case setup from special files (called OpenFOAM dictionaries). For details on the case setup, we refer to the OpenFOAM documentation (<https://cfd.direct/openfoam/user-guide/>, accessed on 9 November 2021). Our simulation case was a cubic channel flow. We used the OpenFOAM blockMesh (<https://cfd.direct/openfoam/user-guide/v9-blockmesh/>, accessed on 9 November 2021) functionality for the mesh creation.

2.3. Molecular-Continuum Coupling

The coupling of the MD model and the continuum model was established using MaMiCo [21,22]. Its purpose is to couple an arbitrary massively parallel MD solver with a continuum solver. The solvers are completely encapsulated. This allows one to exchange one of the solvers without influencing the overall setup. The coupling algorithm is based on the approach by Nie et al. [9]. An overview of the implementation for the coupling with MaMiCo is given below. For further information on MaMiCo and its functionalities, we refer to [21–23].

The overall simulation domain was split into two overlapping parts. One part was assigned to the continuum solver V_C and the other to the MD solver V_M . The intersection of both regions was the overlap V_O . For the data transfer between the two solvers a Cartesian grid-like data structure called *macroscopic cells* was applied. For sake of simplicity, we used the same setup for the macroscopic cells as for the continuum grid; however, other configurations would be possible, while requiring additional interpolation routines.

Within the overlap region, both solvers needed to be physically consistent. A state-flux coupling of the velocity u and the density ρ was applied here: fluxes from the continuum solver were applied to the MD system, whereas the MD state was sampled and imposed as a boundary condition to the continuum solver. Two things were necessary to establish this: first, the transfer of the data between the two solvers, and second, the correct application of the data as a boundary condition in the respective solver. This included a mapping from the molecular to continuum data and vice versa:

Molecular to continuum: The velocity was sampled over all particles, which were located in the region of a macroscopic cell (which is equal with the continuum cell). The sampled velocity was applied as a Dirichlet boundary condition on the corresponding boundary. Since an incompressible continuum solver was used, the density was not manipulated.

Continuum to molecular: An additional force [9] was applied to every particle to match the fluid acceleration from the corresponding continuum cell. The continuity of mass was achieved by either application of periodic boundary conditions or inserting or deleting particles in the outermost macroscopic cells of the MD domain in accordance to the mass flow in the continuum at this boundary.

Besides the spatial dimensions of the coupling, the time integration scheme for the coupling shall also be mentioned here. Therefore, we will sketch the process of one coupled time step in the middle of the simulation. The explanation is supported by Figure 1, which illustrates the cycle. In the initial state, the coupled simulation finished the n -th timestep (t^n). More precisely, the continuum flow field was available at time step n , and MD flow

information was available as sampled data for time point t^n , which were generated from the single MD time points between $t^{n-0.5}$, $t^{n+0.5}$. The simulation of the next coupled time step t^{n+1} works as follows:

1. The continuum solver evolves from time step n to $n + 1$. The boundary condition for this calculation is deduced from the flow field u^n of the molecular solver at time step n . Now the flow field for the continuum part at time step $n + 1$ is available (excluding the boundary values on the MD boundary).
2. The MD solver advances over several MD time steps from continuum time step $n + 0.5$ to $n + 1.5$. The boundary condition for this process is deduced via extrapolation and interpolation of the two latest flow fields of the continuum solver, which are u^n and u^{n+1} . When this is finished, the flow field for the molecular part of the domain is calculated by averaging all time steps from $t^{n+0.5}$ to $t^{n+1.5}$. The corresponding averages in grid cells, which represent boundaries of the continuum solver, are used as boundary condition on the continuum side, and thus complement the continuum flow field data at t^{n+1} .




Figure 1. Time integration and extrapolation scheme of the coupled molecular-continuum simulation from time step n to time step $n + 1$, based on the time staggering of MD and CFD.

2.4. Implementation: Coupling OpenFOAM and MaMiCo

The coupling of OpenFOAM and MaMiCo can be implemented in different ways. One could be to use the OpenFOAM function “externalCoupled” (www.openfoam.com/documentation/guides/latest/doc/guide-fos-field-externalCoupled.html, accessed on 9 November 2021). It couples internal OpenFOAM solvers with external solvers, and it is based on the reading and writing of files. One of MaMiCo’s design principles is, however, the efficient execution of large-scale, concurrently coupled simulations on high performance computing systems. Since concurrently coupled systems require frequent exchange of information, file-based data exchange cannot be part of this coupling. Therefore, the “externalCoupled” tool was not used. Instead, the coupling was designed by the implementation of direct data exchange via a MaMiCo interface to the OpenFOAM solver icoFoam.

The only communication between OpenFOAM and MaMiCo happened through this interface. Thereby, the continuum and the molecular solver were completely encapsulated. The interface is specific for the test case, the Couette startup, we used, but one could adapt it to another test case or to be generic in the future. Due to the small spacial domain, the parallelization of the icoFoam solver would not have been useful. Thus, the interface was implemented for a sequential continuum solver.

The interface for a continuum solver for MaMiCo needs to implement the following independent functionalities:

1. Initialization of the continuum solver and communication of the general setup to the solver (e.g., time step and domain size).
2. Trigger the continuum solver to solve the next time step and update the velocity field.
3. Return the velocity at a certain position in the continuum mesh for the current time step.
4. Prepare the continuum solver to receive and apply the data from the molecular solver as a boundary condition.

5. Receive the values for the current time step from the molecular solver and apply them as boundary conditions within the continuum solver fields.
6. Write output files for the current time step.

These functions enable MaMiCo to control the simulation and the coupling. For the correct initialization of the simulation, functions 1 and 4 is needed. Within one coupled time step, it needs function 2 to advance the continuum solver and functions 3 and 5 to exchange the data between the two solvers. Based on the user's setup, MaMiCo uses function 6 to get the output from the continuum solver.

For the implementation of these functions in the icoFoam interface for MaMiCo, the icoFoam executable needs to be split. The code of the icoFoam script is copied to the corresponding functionalities of the interface. The initialization goes to function 1, and the update for the flow field variables for the next time step is function 2. Some of the interface's functionalities need to be written anew. This is the case for the functions connected with the coupling functions 4 and 5, and also function 6, since the output is given in similar manner to the other MaMiCo output. The getter function for the velocity is basically a wrapper of the OpenFOAM getter for the flow field.

3. Results

We simulated a Couette flow startup with the coupled molecular-continuum model described above. The parametric configurations were varied to identify their effects on the overall quality of the simulation result. We start with a default setup, which built a baseline for our experiments and delivered physically consistent output, close to the analytical solution. To simplify considerations, we restricted ourselves to one-dimensional parameter variations, i.e., varying only one parameter and keeping all other parameters fixed. The change was analyzed by measuring the error $e(t^n)$ in time step t^n , which means the difference between the analytical and the simulated solution for every grid cell:

$$e(t^n) = \sqrt{\frac{1}{C} \sum_{\vec{c}_j, \vec{c}_j \notin V_O} |u_s(t^n, \vec{c}_j) - u_a(t^n, \vec{c}_j)|^2} \quad (2)$$

where \vec{c}_j denotes the center point of the j -th grid cell; $u_s(t^n, \vec{c}_j)$ is the simulated velocity from the molecular dynamics in x -direction at \vec{c}_j in time step t^n ; $u_a(t^n, \vec{c}_j)$ is the analytical solution for the flow velocity; and C is the number of grid cells considered. This yielded the specific error for every time step. We further quantified the mean error over time from an average over all MD time steps N , apart from the 100 first coupling time steps, to exclude initial non-equilibrium perturbations from switching on the coupling:

$$e = \frac{1}{N} \sum_{n=100}^N e(t^n). \quad (3)$$

In the following, we list the technical parameters we analyzed.

Flow velocity magnitude: The analytical solution of the continuum model is only correct for an incompressible fluid. This assumption can be made for fluids with small Mach numbers $Ma < 0.3$ [24]. The Mach number is the relation of the flow speed to the speed of sound. Increasing the overall flow velocity will increase the Mach number and thereby amplify compressibility effects and related modelling error.

Cell size of the continuum solver: The result from the molecular data needs to be sampled before it can be applied in the continuum simulation. To calculate a value, for example, the velocity for a continuum cell, the mean velocity of all particles within the area of this cell is analyzed. Under the condition of a constant density, the number of particles included in the calculation depends on the size of the continuum cell: the bigger the cell, the more particles used for sampling, and the lower the thermal fluctuations; see, e.g., [25] or the discussion on "Number of MD instances in a multi-instance sampling" below.

Coupling time step: One coupling time step includes one time step in the continuum and several time steps in the MD simulation, due to the smaller time scale in MD. Increasing the coupling time step results in more MD time steps for sampling, and thus, in smoother data. Making the coupling time step too large, however, also impacts the continuum solver's stability. In the worst case, this yields oscillations of up to an entire divergence of the coupled simulation. Nie et al. discussed the limits for the time step [9]. The upper bound is defined by the characteristic time of the flow and the continuum grid. The lower bound is given by the fact that the time step needs to be larger compared to the velocity auto-correlation time, which is 0.14τ for a Lennard–Jones fluid with density $\rho = 0.81$ and temperature $T = 1.1$ [9], where τ is the characteristic time of the Lennard–Jones potential. Both together yield the following relation, using the spacial discretization Δx , the time step Δt and the kinematic viscosity of the fluid ν :

$$0.14\tau < \Delta t < \frac{\Delta x^2}{\nu} \quad (4)$$

which in our dimensionless case results in $\Delta t \in [0.14, 9.47]$.

Number of MD instances in a multi-instance sampling: Multi-instance sampling refers to sampling over an ensemble of quasi-identical, quasi-independent, randomized MD simulations. Due to its Monte Carlo-like nature, the fractional error E_u of velocity scales with [25]

$$E_u \propto \frac{1}{\sqrt{M}}$$

where M is the number of MD instances. This means, doubling the number of instances reduces the error only by a factor of 0.7.

Region for the momentum insertion in the MD model: Both the location and the thickness of the overlap region, in which momentum is transferred from continuum to MD, impact the overall quality of the simulation solution. For example, when using two cell layers instead of one, the MD simulation implicitly also obtains information on the velocity gradient from the continuum simulation, which typically leads to better agreement of the two solvers. Wang et al. found that better results were obtained in their coupled molecular-continuum model, when the area for continuum-to-particle transfer was increased [16].

Setup for the mass insertion in the MD model: The simulation of a flow over the boundary of the MD domain was accomplished by insertion and deletion of mass. From physical and mathematical point of view, this algorithmic step has a severe impact: mathematically, adding/removing a molecule resembles the incorporation/removal of ordinary differential equations; physically, the non-linear energy landscape must not be perturbed severely [26], and so forth.

Table 1 summarizes the simulations we performed. The list includes the special parameters analyzed and the error within the MD domain calculated, as described in Equation (3). The other, general parameters and details for the coupled simulations, which are not listed, can be found in the Appendix A.

For the default simulation, a “good” setup was chosen, exhibiting a small statistical error of 0.007. Exemplary parts of the velocity field of this simulation are shown in Figure 2a. As the small error suggests, the velocities for the simulation were in good agreement with the analytical solution during all time steps. Only minor fluctuations from MD can be seen to slightly influence the continuum. Stability of the overall scheme was preserved.

For simulation 1, the number of MD instances was reduced from 200 to 50. Due to this, the fluctuations within the MD simulation were increased, as can be seen in Figure 2b. Again, the fluctuations from the molecular dynamics domain had an influence on the velocity in the cells of the continuum that were close to the boundary. Regardless of this effect, the coupled simulation remained stable, and the correct steady state was reached. The error increased by a factor of two compared with the default simulation, which is in line with kinetic theory.

Table 1. Setup of the technical parameters for the different simulations and error of the simulated velocity compared to the analytical solution within the molecular domain. The focus is on the parameters, which were varied and analyzed.


Name	MD Instances	Coupling Time Step	Wall Velocity in x -Direction	Momentum Insertion	Boundary Condition	Particle Insertion/Removal	Error MD e
default	200	0.25	0.5	2–3	periodic	none	0.007
sim-1	50	0.25	0.5	2–3	periodic	none	0.014
sim-2	200	1.0	0.5	2–3	periodic	none	0.015
sim-3	200	0.25	1.5	2–3	periodic	none	0.015
sim-4	200	0.25	0.5	1	periodic	none	0.007
sim-5	200	0.25	0.5	2	periodic	none	0.011
sim-6	200	0.25	0.5	3	periodic	none	0.008
sim-7	200	0.25	0.5	2–3	reflecting	USHER	0.010

The next parameter under consideration was the coupling time step, which was increased by a factor 4 to 1.0 for simulation 2. The effect of this variation can be seen in Figure 2c. At first glance, the velocity field seems to be in good agreement with the analytical solution. Compared to the default simulation, where the difference was random because of the fluctuating velocity, the velocity in this case was systematically slower than it should have been, compared to the analytical velocity. Due to the longer time step, the temporal evolution of the sudden startup of the Couette flow could not be resolved sufficiently; consequently, this continuum phenomenon also manifested in the coupled case. Close to the stationary state, in the analytical in simulation results the difference still exists. Velocities from both the continuum and the MD solver did not recover from the initial time step size-induced error.


In simulation 3, we analyzed the influence of a higher overall velocity. Therefore, we increased the velocity of the moving wall in x -direction from 0.5 to 1.5. The resulting velocity field is plotted in Figure 2d. Due to the higher velocity, which did not influence the statistical fluctuations of the molecules, the signal-to-noise ratio, as defined by Hadjiconstantinou, was enlarged compared to the baseline case [25]. This change gives the impression that the agreement of the simulated and analytical solutions improved, even though the quantitative error within the domain of the MD was higher, at 0.015, compared to the simulation with a slower velocity. Due to the higher velocity magnitude, the Mach number $Ma = 0.26$ was three times higher compared to the baseline scenario, which is still slightly below the usual bound $Ma < 0.3$ that is used to justify the assumption of incompressibility. Nevertheless, the observed error must be attributed to the influence of compressibility.

Next, we considered the size of the overlap region for momentum insertion. Until this case, two cells were used. This was reduced to one cell, and the position of this cell layer was varied. We started by only inserting momentum in the outermost cell layer of the overlap region in terms of the MD domain (simulation 4). The velocity profile is not shown, due to its similarity to the ones from the default case. Additionally, the error in the MD region is of comparable size for both setups.


In simulation 5, only one cell layer was used for momentum transfer, as in simulation 4. This time, however, this cell layer was placed in the middle of the overlap region, which means in the second layer of overlapping cells. The simulated velocity for this situation can be found in Figure 2e. In contrast to simulation 4, this setup differed from the baseline case. A lag in the profile, compared to the analytical solution, could be already observed for the 404-th coupling time step. This lag stayed until the 1304-th time step, and from there on, slowly decreased, but remained until the end of the simulation.




(a) default




(b) sim-1: 50 instead of 200 instances, compared to default




(c) sim-2: time step 1.0 instead of 0.25, compared to default



(d) sim-3: wall velocity 1.5 instead of 0.5, compared to default



(e) sim-5: momentum insertion in second cell layer



(f) sim-7: run with mass insertion and removal

Figure 2. Simulation results: velocity u in x -direction shown over the z -coordinate. The lines refer to the analytical solutions for the Couette flow startup; the circles show u from the continuum model, and the crosses u for the MD model. The colors stand for coupling time steps (e.g., the legend entry “Step 104” refers to coupling time step 104).

By moving the momentum insertion from the outermost to the middle cell layer of the overlap domain, the error increased. We therefore shifted the one-cell-momentum-transfer region by one more cell layer into the third overlap layer, to better understand this effect. This setup was chosen for simulation 6. The results indicate that the error decreased. The error in this simulation was 0.008; i.e., this setup was slightly worse than the baseline case and simulation 4. However, again the analytical and the simulated velocity profiles are in good agreement (not shown in Figure 2). To conclude regarding the last three simulations, decreasing the overlap layer to, e.g., only one cell, increased the probability of error and resulted in higher sensitivity of the overall coupling.

Next, we analyzed the impact of mass insertion and deletion at all outer boundaries of the MD domain. The details on the insertion and deletion of mass can be found in the Appendix A together with all the details of the simulation setup. To emphasize the impact of particle insertion and removal, we further modified the MD boundary condition and changed the periodic boundaries into reflecting ones. This setup is referred to as simulation 7, and velocity profiles can be found in Figure 2f. In the simulation with these parameters, the error in the velocity field was around 0.010, close to the accuracy of the default setup (without mass transfer). As stated before, a small loss in accuracy was expected. In simulation 7, the insertion or deletion of particles modifies the mass in the system, and therefore, the validity of the density requires analysis. The scenario was initialized with a density of 0.8130. In the default simulation, the density after 3000 coupling cycles was 0.8115 in the inner part of the MD domain, which corresponds to a deviation of ca. 0.1%, due to minor perturbations at the outer boundaries evolving from the boundary forcing. The error in the density field of the simulation 7 increased by a factor of 5 to 0.49%, but still remained small. The particle distribution within the molecular dynamics domain was homogeneous. Due to the periodic boundary condition for the MD domain in the default simulation, density remained constant and no discrepancies were observed along the channel.

Finally, we analyzed the temporal evolution of the error and related fluctuations within the molecular dynamics domain by comparing of all eight setups, as shown in Figure 3, where we plotted $e(t^n)$ (see Formula (2)). The default case showed an almost constant error over time. Simulations 4 and 6, which only used the outermost or innermost cell for the momentum insertion, had slightly increased error during their respective initial phases. From $t = 300$ onward, the error rates of simulations 4 and 6 were comparable to that in the default case.




Figure 3. The temporal evolution of e for the different setups.

As expected, the setup using less instances for the multi-instance sampling was similar to the baseline case in the temporal evolution by showing a constant, yet higher error rate.

In the other four setups, we observed an initial increase in the error, which after a while decreased again, and in some cases, stabilized, though still exceeding the error of the default case. The higher error rates at the beginning agree with our analyses of Couette profiles. We also conducted a deeper analysis of the thermal fluctuations in our Couette flows with varying MD ensemble sizes, which were generally in good agreement with statistical mechanics; refer to [27] for details.

Furthermore, we wanted to comment on the influence of the continuum cell size by increasing it by a factor of 1.5. For this purpose, we increased the entire setup size (cf. Table 2), and positioned the (lower, left, front corner of the) MD domain at an offset (30,30,15). All other parameters stayed constant and can be found in the Appendix A.

Table 2. Parameters and the error within the MD for the second set of simulations. The focus is on the parameters which were varied.

Name	MD Instances	MD Domain Size	Channel Width	Continuum Cell Size	e
sim-8	200	90, 90, 90	150	5.0, 5.0, 5.0	0.0075
sim-9	59	90, 90, 90	150	7.5, 7.5, 7.5	0.0073

To match the overall number of samples in simulations 8 and 9 [25], we ran the simulations at resolutions 5.0 and 7.5, respectively, with modified MD ensembles. Both the simulated velocity profiles had good agreement with the analytical solution (not shown here). The setup sim-9 exhibited a slightly smaller (approximately 3%) but comparable error.

4. Discussion and Lessons Learned

In this contribution, we analyzed hybrid transient molecular-continuum simulations with various setups of technical parameters. The comparison was based mostly on the quantitative and qualitative differences in the velocity field.

Based on our observations reported in the previous sections and all the work we conducted beyond, we formulated some key lessons learned for establishing a stable and accurate two-way coupled molecular-continuum flow simulation. Our focus here was put on Couette flow, which is a one-dimensional flow scenario, but it still allowed us to draw conclusions even for fully three-dimensional systems. In the future, we aim to apply this molecular-continuum simulation technology to systems that focus on relevant scenarios from process engineering and thermodynamics, such as evaporation, or multiphase scenarios.

4.1. Lesson 1: Averaging Matters!

Despite the usual literature in the field, which often states that the coupling direction molecular→continuum is “straight-forward” via simple averaging, we want to stress few key observations. Statistical error analysis tells us that (a) independent samples and (b) a sufficiently big set of samples are required, which easily yields very computationally intensive setups—computationally intensive in the sense of spatial, temporal or ensemble averaging (multi-instance sampling). While 1D-flows such as Couette flow can be tackled through sampling over the other two dimensions in a 3D setup, our default simulation setup was chosen under the implicit requirement of a fully 3D flow scenario. I.e., with our default simulation and simulation 1, we could get a rough answer on how many samples we actually required, which were not fully independent, as we sampled over all time steps. Nevertheless, considering the given flow scenario (flow velocities, Mach number, etc.) with density $\rho = 0.81$, macroscopic cells of size $dx = 5.0$ and 50 MD time steps per coupling time step, our multi-instance approach yielded approximately 100 million/25 million molecule samples in the default simulation/simulation 1, per cell. This justifies (a) the use of high-performance and supercomputing for general molecular-continuum setups and

(b) the need for more enhanced algorithms to reduce Brownian noise, such as algorithms incorporating noise filters; see also, amongst others, [23,28].

All these points are mentioned here with regard to a coupling with an incompressible or a weakly compressible solver. Based on other research [4], these challenges are expected to be less pronounced for fully compressible systems.

While providing significant benefits in terms of reducing the computational effort, noise filters, however, introduce even more parameters to tune. This renders actual a priori error estimation more complex. More work in this direction is necessary and also under current investigation by our group.

4.2. Lesson 2: Only Long-Term Simulation Runs Tell the Truth!

As observed in, e.g., our study on the location and thickness of the overlap layer, actual stability and accuracy cannot necessarily be detected from short-term simulations. Artifacts can manifest themselves only slightly in the beginning but become more pronounced in the long-term. Therefore, long-term runs are advisable to numerically demonstrate the stability and accuracy of a molecular-continuum method.

4.3. Lesson 3: Choose the Overlap Layer Big Enough!

While reported in the literature before (see, e.g., [16]), we confirmed with our studies that the size and shape of the overlap region has a severe impact on the molecular-continuum simulation quality.

4.4. Lesson 4: Continuum Scales Matter!

As mentioned in Lesson 1, the number of samples required for stable and accurate coupling can easily grow to millions when using molecules. If ultra-fine resolution is also required on the continuum side, e.g., to resolve geometries, this will become an issue, even if the simulation scenario allows for a longer coupling time step—which can easily not be the case due to accuracy or stability limits of the continuum solver, e.g., due to strong gradients in the flow (cf. simulation 2).

To meet sampling requirements, we considered macroscopic cells of size $5 \times 5 \times 5$, making them bigger than the cells used in our prior work [22]. Our multi-instance sampling approach is an additional ingredient for coping with this challenge.

4.5. Lesson 5: Modular Programming, Testing, Interoperability and Portability Are Essential to Robust Coupling!

Although this is a well-known fact for most complex computational problems in science and engineering [29], we want to stress that software design, modular programming and (automated) unit and integration testing are essential requirements for long-term simulation software development. Indeed, we were faced with the challenge of fixing a few software bugs that were hard to detect and only became visible in large-scale/long-term runs through minor perturbations in the overall solution. Our software MaMiCo already incorporates several of these kinds of tests, and we are currently working towards extending upon this in the near future.

Similarly, interoperability—in our case: interfacing various MD and continuum solver software packages—is a major challenge, and in this contribution, we outlined how to also include bigger continuum solver software packages with the example of OpenFOAM.

Portability is also important, not only for porting software from laptops to supercomputers, but also for porting between different supercomputing platforms. Amongst others, we currently use the ARM A64FX platform for molecular dynamics simulations, which requires extra instructions (e.g., SVE) to optimally exploit the hardware. Although it is well-known that MD is the most computationally intensive part in a molecular-continuum system, portability aspects are also relevant for the interfacing and coupling software, such as MaMiCo.

Author Contributions: Conceptualization, H.W. and P.N.; methodology, H.W.; software, H.W.; validation, H.W. and P.N.; data curation, H.W.; writing—original draft preparation, H.W.; writing—review and editing, P.N.; visualization, H.W.; supervision, P.N.; project administration, P.N.; funding acquisition, P.N. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially funded and supported by the project “MaST: Macro/Micro-Simulation of Phase Separation in the Transcritical Regime” of the center for digitization and technology research of the armed forces (dtec.bw), and the HSU-internal research funding (IFF) project “Resilience and Dynamic Noise Reduction at Exascale for Multiscale Simulation Coupling”.

Data Availability Statement: The analysed results can be reproduced. Therefore, everything may be found in the *SensitivityRuns* branch of the MaMiCo Github repository (<https://github.com/HSU-HPC/MaMiCo/tree/SensitivityRuns/coupling>, accessed on 9 November 2021).

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

CFD	Computational fluid dynamics
MD	Molecular dynamics
C	number of grid cells
\vec{c}_j	center point of the j -th grid cell
dx	macroscopic cell size
E_u	the fractional error of the velocity
e	measured error
\vec{F}_i	force acting on molecule i
H	width of the channel
M	number of independent MD instances
Ma	Mach number
m	mass of a molecule
N	total number simulated MD time steps
T	temperature
t	time
t^n	the n -th time step of the simulation
u_a	analytical calculated velocity
u_{CFD}	velocity in x -direction calculated by the CFD
u_{MD}	velocity in x -direction calculated by the MD
u_s	simulated velocity
u_w	velocity at the moving wall of the Couette scenario
V_C	domain of the continuum solver
V_M	domain of the molecular solver
V_O	overlap domain of the two solvers
\vec{v}_i	velocity of molecule i
\vec{x}_i	position of molecule i
Δt	time step
Δx	spacial discretization
ϵ	characteristic energy parameter of the Lennard–Jones potential
ν	kinematic viscosity of the fluid
ρ	density
σ	characteristic length scale for the MD
τ	characteristic time of the Lennard–Jones potential

Appendix A. Details of the Simulation Setup

Appendix A.1. Couette Flow Startup

A 3D channel flow is considered. The wall at $z = 0$ is moving with a constant velocity u_w in x -direction. The default velocity for the wall is set to 0.5. The fluid is initialized with

a zero velocity in the whole flow field. It is subsequently accelerated by the moving wall. The analytical solution for the x -velocity u in dependence of the position in z at time t is

$$u(z, t) = u_w \left(1 - \frac{z}{H}\right) - \frac{2u_w}{\pi} \sum_{k=1}^{\infty} \frac{1}{k} \sin\left(k\pi \frac{z}{H}\right) \exp\left(-k^2 \pi^2 \frac{vt}{H^2}\right) \quad (\text{A1})$$

The formula includes the width of the channel H . In our analysis, we use an approximation which includes all terms up to $k = 29$.

Appendix A.2. Setup for the Molecular Dynamics Simulation

The molecular dynamics simulation is conducted using SimpleMD [21]. The domain is cubic with a size of $60 \times 60 \times 60$ and periodic boundary conditions in x - and y -direction. In the z -direction, reflecting boundary conditions are applied together with the Zhou boundary force to emulate the missing particles beyond the boundary [19]. In some study cases, the USHER algorithm [26] is applied at the boundary in combination with reflecting boundary conditions and the Zhou boundary force. The MD time step is chosen to be 0.005. The cutoff radius is set to 2.2. A thermostat is used to apply a constant temperature of 1.1 within each macroscopic cell.

At the beginning of a simulation, $56 \times 56 \times 56$ particles are evenly distributed along all three spatial directions, the resulting density is approximately 0.81. The simulation is equilibrated for 10,000 time steps before the coupled simulation is started. Multi-instance sampling is applied for the coupled simulations. This means, the output from the molecular dynamics simulations is the averaged data from several independent molecular dynamics simulations. In the default setup, we use 200 instances for the sampling.

Appendix A.3. Setup for the Continuum Solver

The icoFoam solver from the OpenFOAM library is used. The domain as well as the grid are cubic, where the former has a size of $100 \times 100 \times 100$ and the latter $5 \times 5 \times 5$. The cube has a missing cubic part in the middle due to the existence of the MD domain, which has an offset compared to the CFD coordinates of (20,20,5). At this inner boundary, a Dirichlet boundary condition is applied for the velocity, according to the values from MD. A Neumann boundary condition is applied at every boundary for the pressure and at the outer boundary for the velocity. The time step is set to 0.25. The density is set to 0.81 and the kinematic viscosity 2.64, which is consistent with the MD model.

Appendix A.4. Dimensions in the Simulations

All variables used for the solvers are given in dimensionless form. Corresponding to the reduced Lennard–Jones units, the variables are reduced using the characteristic length σ , the energy parameter ϵ and the mass of one particle m .

References

- O'Connell, S.T.; Thompson, P.A. Molecular dynamics–continuum hybrid computations: A tool for studying complex fluid flows. *Phys. Rev. E* **1995**, *52*, R5792–R5795. [[CrossRef](#)] [[PubMed](#)]
- Barsky, S.; Delgado-Buscalioni, R.; Coveney, P.V. Comparison of molecular dynamics with hybrid continuum–molecular dynamics for a single tethered polymer in a solvent. *J. Chem. Phys.* **2004**, *121*, 2403–2411. [[CrossRef](#)] [[PubMed](#)]
- Borg, M.; Lockerby, D.; Reese, J. Fluid simulations with atomistic resolution: A hybrid multiscale method with field-wise coupling. *J. Comput. Phys.* **2013**, *255*, 149–165. [[CrossRef](#)]
- Delgado-Buscalioni, R.; Flekkøy, E.; Coveney, P. Fluctuations and continuity in particle-continuum hybrid simulations of unsteady flows based on flux-exchange. *EPL (Europhys. Lett.)* **2005**, *69*, 959–965. [[CrossRef](#)]
- Dupuis, A.; Kotsalis, E.M.; Koumoutsakos, P. Coupling lattice Boltzmann and molecular dynamics models for dense fluids. *Phys. Rev. E* **2007**, *75*, 046704. [[CrossRef](#)]
- Flekkøy, E.; Wagner, G.; Feder, J. Hybrid model for combined particle and continuum dynamics. *EPL (Europhys. Lett.)* **2007**, *52*, 271. [[CrossRef](#)]
- Kalweit, M.; Drikakis, D. Coupling strategies for hybrid molecular-continuum simulation methods. *Proc. Inst. Mech. Eng. Part C J. Mech. Eng. Sci.* **2008**, *222*, 797–806. [[CrossRef](#)]

8. Li, Y.; Liao, J.; Yip, S. Nearly Exact Solution for Coupled Continuum/MD Fluid Simulation. *J. Comput.-Aided Mater. Des.* **2001**, *6*, [CrossRef]
9. Nie, X.B.; Chen, S.Y.; E, W.N.; Robbins, M.O. A continuum and molecular dynamics hybrid method for micro- and nano-fluid flow. *J. Fluid Mech.* **2004**, *500*, 55–64. [CrossRef]
10. Sun, J.; He, Y.; Tao, W. Scale effect on flow and thermal boundaries in micro-/nano-channel flow using molecular dynamics–continuum hybrid simulation method. *Int. J. Numer. Methods Eng.* **2010**, *81*, 207–228. [CrossRef]
11. Wagner, G.; Flekkøy, E.; Feder, J.; Jøssang, T. Coupling molecular dynamics and continuum dynamics. *Comput. Phys. Commun.* **2002**, *147*, 670–673. [CrossRef]
12. Werder, T.; Walther, J.H.; Asikainen, J.; Koumoutsakos, P. Continuum-particle hybrid methods for dense fluids. In *Multiscale Modelling and Simulation*; Lecture Notes in Computational Science and Engineering; Springer: Berlin/Heidelberg, Germany, 2004; pp. 227–235. [CrossRef]
13. Mohamed, K.M.; Mohamad, A.A. A review of the development of hybrid atomistic–continuum methods for dense fluids. *Microfluid. Nanofluid.* **2009**, *8*, 283–302. [CrossRef]
14. Ren, W. Analytical and numerical study of coupled atomistic-continuum methods for fluids. *J. Comput. Phys.* **2007**, *227*, 1353–1371. [CrossRef]
15. Yen, T.H.; Soong, C.Y.; Tzeng, P.Y. Hybrid molecular dynamics-continuum simulation for nano/mesoscale channel flows. *Microfluid. Nanofluid.* **2007**, *3*, 729. [CrossRef]
16. Wang, Q.; Ren, X.G.; Xu, X.H.; Li, C.; Ji, H.Y.; Yang, X.J. Coupling Strategies Investigation of Hybrid Atomistic-Continuum Method Based on State Variable Coupling. *Hindawi-Adv. Mater. Sci. Eng.* **2017**, *2017*, 1014636. [CrossRef]
17. Ren, X.; Wang, Q.; Xu, L.; Yang, W.J.; Xu, X. HACPar: An efficient parallel multiscale framework for hybrid atomistic–continuum simulation at the micro- and nanoscale. *Adv. Mech. Eng.* **2017**, *9*, 168781401771473. [CrossRef]
18. Longshaw, S.; Borg, M.; Ramisetty, S.; Zhang, J.; Lockerby, D.; Emerson, D.; Reese, J. mdFoam+: Advanced Molecular Dynamics in OpenFOAM. *Comput. Phys. Commun.* **2018**, *224*, 1–21. [CrossRef]
19. Zhou, W.J.; Luan, H.B.; He, Y.L.; Sun, J.; Tao, W.Q. A study on boundary force model used in multiscale simulations with non-periodic boundary condition. *Microfluid. Nanofluid.* **2014**, *16*, 587–595; Erratum in **2016**, *20*, 93. [CrossRef]
20. Griebel, M.; Knapek, S.; Zumbusch, G. *Numerical Simulation in Molecular Dynamics—Numerics, Algorithms, Parallelization, Applications*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2007.
21. Neumann, P.; Flohr, H.; Jarmatz, P.; Tchipev, N.; Bungartz, H.J. MaMiCo: Software design for parallel molecular-continuum flow simulations. *Comput. Phys. Commun.* **2015**, *200*, 324–335 [CrossRef]
22. Neumann, P.; Bian, X. MaMiCo: Transient multi-instance molecular-continuum flow simulation on supercomputers. *Comput. Phys. Commun.* **2017**, *220*, 390–402. [CrossRef]
23. Jarmatz, P.; Neumann, P. MaMiCo: Parallel Noise Reduction for Multi-instance Molecular-Continuum Flow Simulation. In *Computational Science—ICCS 2019*; Rodrigues, J.M.F., Cardoso, P.J.S., Monteiro, J., Lam, R., Krzhizhanovskaya, V.V., Lees, M.H., Dongarra, J.J., Sloot, P.M., Eds.; Springer International Publishing: Cham, Switzerland, 2019; pp. 451–464.
24. Ferziger, J.H.; Perić, M.; Street, R.L. *Computational Methods for Fluid Dynamics*, 4th ed.; Springer: Berlin/Heidelberg, Germany, 2019.
25. Hadjiconstantinou, N.; Garcia, A.L.; Bazant, M.; He, G. Statistical error in particle simulations of hydrodynamic phenomena. *J. Comput. Phys.* **2003**, *187*, 274–297. [CrossRef]
26. Delgado-Buscalioni, R.; Coveney, P.V. USHER: An algorithm for particle insertion in dense fluids. *J. Chem. Phys.* **2003**, *119*, 978–987. [CrossRef]
27. Jafari, V.; Wittmer, N.; Neumann, P. Massively Parallel Molecular-Continuum Flow Simulation with Error Control and Dynamic Ensemble Handling. In *Proceedings of HPC Asia 2022*; ACM: New York, NY, USA, 2022; accepted.
28. Zimoń, M.J.; Prosser, R.; Emerson, D.R.; Borg, M.K.; Bray, D.J.; Grinberg, L.; Reese, J.M. An evaluation of noise reduction algorithms for particle-based fluid simulations in multi-scale applications. *J. Comput. Phys.* **2016**, *325*, 380–394. [CrossRef]
29. Råde, U.; Willcox, K.; Curfman McInnes, L.; De Sterck, H. Research and Education in Computational Science and Engineering. *SIAM Rev.* **2018**, *60*, 707–754. [CrossRef] [PubMed]