Connecting Molecular Dynamics and Computational Fluid Dynamics

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus Prof. Ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties
in het openbaar te verdedigen op maandag 16 mei 2011 om 10:00 uur

door **Anton Pieter Markesteijn**

werktuigbouwkundig ingenieur,
geboren te Papendrecht, Nederland

Dit proefschrift is goedgekeurd door de promotoren:
Prof. dr. ir. J. Westerweel
Prof. Dr. rer. nat. S. Luding

Samenstelling promotiecommissie:

Rector Magnificus,	Rector Magnificus, voorzitter
Prof. dr. ir. J. Westerweel,	Technische Universiteit Delft, promotor
Prof. Dr. rer. nat. S. Luding,	Universiteit Twente, promotor
Prof. dr. A.A. Darhuber,	Technische Universiteit Eindhoven
Prof. dr. ir. T.J.H. Vlugt,	Technische Universiteit Delft
Prof. dr. H. Schiessel,	Universiteit Leiden
Prof. dr. ir. D.J. Rixen,	Technische Universiteit Delft
Dr. H. Wijshoff,	Océ NV

ISBN/EAN 9789461081742

Printed by: Gildeprint Drukkerijen - Enschede

Samenvatting

Koppeling van Moleculaire Dynamica en Numerieke Stromingsleer

Eén van de belangrijkste ontwikkelingen van de afgelopen eeuwen is de evolutie van de miniaturisering en het inzicht wat daaruit is voorgebracht. Voor the fysische wetenschappen betekende dit beter begrip van wat materie eigenlijk is, terwijl het vanuit een praktisch oog punt betekende dat er meer geavanceerde toepassingen werden ontwikkeld. De succesvolle ontwikkeling van de miniaturisering vereist zorgvuldige planning, en computer simulaties helpen met dit aspect. Echter, verdere miniaturisering leidt ook tot problemen. Het is algemeen bekend dat de wereld rond ons uit atomen bestaat, nochtans voor algemene macroscopische verschijnselen werkt het concept van een continuüm erg goed. In plaats van afzonderlijke atomen te bestuderen, wordt het gedrag van een zeer groot aantal atomen bestudeerd en wordt deze uitgedrukt in wiskundige vergelijkingen die opgelost kunnen worden met een bepaalde methode. Echter, de afgelopen decennia is er veel interesse getoond in nanotechnologie, d.w.z. de studie van fenomenen die gebeuren op de nano-schaal (< 0.0001mm) waar de materie wordt gemanipuleerd op de atomische en moleculaire schaal. Het concept van een continuüm zal in deze gevallen falen en andere technieken moeten worden gebruikt om deze verschijnselen te simuleren. Er bestaat echter ook nog een tussenliggend gebied (waar de typische grootte ~ 0.00001mm-~ 0.01mm), waar zowel de moleculaire effecten en de continuüm effecten van belang kunnen zijn. Om dit te simuleren met één enkele methode is zeer uitdagend. Bijvoorbeeld, een methode die zeer nauwkeurig op de kleine schaal is, zal al snel te moeizaam zijn om op te lossen op de grote schaal. Aan de andere kant, methoden die over het algemeen efficiënt zijn voor de grote schaal, hebben alle details op de kleine schaal verloren. Een oplossing voor dit probleem is om alleen die specifieke methode in een gebied te gebruiken waar deze het meest geschikt is en zo bij elkaar het gehele domain op te lossen. Echter, om dit mogelijk te maken, zullen beide methoden met elkaar moeten communiceren. In de praktijk betekent dit, dat de methoden worden gekoppeld. Het onderwerp van dit proefschrift is het ontwikkelen, implementeren en testen van één van deze methoden, die algemeen bekent staan als multiscale, gekoppelde, of hybride methoden.

In het huidige onderzoek, is de zogenaamde "Schwarz alternating method" gekozen om een domein dat een vloeistof simuleert met behulp van moleculaire dynamica (MD) en een domein dat continuüm methoden gebruikt, te koppelen. De methode koppelt de twee domeinen op het beginsel dat de twee domeinen dezelfde oplossing hebben in een regio waar ze elkaar overlappen. Binnen dit overlappinggebied wisselen de twee domeinen randvoorwaarden uit met elkaar. De randvoorwaarden voor het continuüm domein kunnen gemakkelijk worden verkregen uit de MD simulatie. De specificatie van randvoorwaarden voor het MD domein is minder eenvoudig. In de huidige implementatie moeten drie belangrijke stappen worden opgelost, welke worden besproken in sectie 3.2. Dit is gedaan voor zowel de koppeling van vloeistoffen die bestaan uit één atoom, zoals argon en meer complexe vloeistoffen, zoals water. Voordat de koppeling kan worden gerealiseerd, is het zinvol om eerst te kijken waar en wanneer er een koppeling mogelijk is en onder welke voorwaarden deze simulaties een juiste oplossing geven. Daarom wordt in het huidige onderzoek ook door middel van MD simulaties/studies onderzocht wat de mogelijkheden en beperkingen zijn van pure MD simulaties, terwijl hetzelfde wordt gedaan voor puur continuüm methoden.

Uit de resultaten van de MD simulaties is gebleken dat grote afwijkingen tussen continuümmechanica en MD vooral merkbaar zijn in de buurt van de wanden van nanokanalen of in de buurt van nano-obstakels. Deze afwijkingen, zichtbaar als grote variaties in de verzamelde

(continuüm) macroscopische variabelen uit de MD simulatie, zijn het resultaat van de interactie van de atomen in de vloeistof met atomen in de wand, welke ook experimenteel aangetoond kunnen worden. Verder weg van de wand tonen de macroscopische variabelen de (verwachte) continuüm waarde. Echter, hoewel deze verschillen aantonen dat er grote niet-continuüm effecten plaatsvinden, is het zo dat zelfs voor kleine nano kanalen van ongeveer 5 nm, continuüm-achtig gedrag kan worden waargenomen. Dit feit werd gebruikt om de viscositeit als functie van de temperatuur voor vier verschillende water modellen te bepalen. In het algemeen bleek dat in een kanaal met een hoogte van ongeveer 8 nm, kon worden gesproken over een continuüm. Echter, dit betekent nog niet dat een puur continuüm methode voor het berekenen van de stroming in deze kanalen aan te raden is. Er zijn ook andere redenen om niet een continuüm methode te gebruiken, want MD simulaties hebben sommige unieke voordelen, zoals de mogelijkheid om realistische wand-vloeistof interacties te voorspellen. Bovendien, met MD kunnen verschillende andere verschijnselen worden gesimuleerd die moeilijk of zelfs onmogelijk zijn met een continuüm techniek, zoals het afbreken van een nano-jet. Het is echter gebleken dat veelgebruikte waarden om de moleculaire interacties af te breken te laag zijn voor het nauwkeurig modelleren van enkele belangrijke fenomenen. Aan de andere kant, bij de vergelijking van continuüm met experimentele resultaten bleek dat voornamelijk de continuüm techniek die elektrokinetische effecten beschrijft, nauwkeurig genoeg is voor zelfs een toepassing waar de hoogte slechts 150 nm is.

De resultaten van de gekoppelde simulaties worden getoond in de sectie 3.3. Hier wordt uitgelegd hoe een gekoppelde simulatie in principe kan worden gezien als een methode om nieuwe randvoorwaarde voor een continuüm domein te verkrijgen. Hier wordt de waarde nu nauwkeuriger bepaald door de communicatie tussen het MD en continuüm domein. Dit werd aangetoond door Poiseuille stroming van argon en water in een groot kanaal te simuleren. De niet-continuüm verschijnselen in de buurt van de wand worden nauwkeurig gesimuleerd door MD terwijl er geen lange MD rekentijd wordt verspild in het gedeelte dat zich gedraagt als een continuüm. De koppeling van MD en continuüm kan ook worden gebruikt om niet-periodieke randvoorwaarden voor MD systemen te specificeren, welke moeilijk of onmogelijk zijn om te implementeren in een pure MD simulatie. Het principe wordt aangetoond door middel van een tweedimensionale gekoppelde simulatie van een "nanowire" in een uniforme stroming van argon. De belangrijkste voordelen en de resultaten van dit soort simulaties zijn dat de invloed op de stroming van één obstakel wordt onderzocht in plaats van de één en al zijn periodieke beelden. De manier hoe de domeinen gekoppeld worden, kan ook worden toegepast voor andere macroscopische variabelen dan snelheid, bijvoorbeeld temperatuur. In dit onderzoek werd een kwalitatieve studie uitgevoerd waar een voorwerp in een temperatuurgradiënt is geplaatst door middel van de koppeling om thermophorese in vloeistoffen te onderzoeken.

Tenslotte, in hoofdstuk 4 wordt een ander soort koppeling tussen moleculen en het continuüm uitgelegd die zeer efficiënt is om het gedrag van polymeren te bestuderen. Voor dit doel, wordt een meso-schaal simulatie uitgevoerd voor de bepaling van de sterkte van de stroming die nodig is een polymeer in een smal kanaal te duwen. De resultaten laten een zeer goede overeenstemming zien met een voorspelling gebaseerd op een de Gennes blob model van het polymeer, welke stelt dat de kritische snelheid flux voor translocatie lineair afhankelijk is van de temperatuur, maar onafhankelijk is van de lengte van de polymeerketen of de breedte van het kanaal.

Summary

Connecting Molecular Dynamics and Computational Fluid Dynamics

One of the most important developments in the last centuries is the process of miniaturisation and understanding everything that it entails. For the physical sciences this meant the continuing discovery what matter is and how it behaves, while from a practical point of view it meant that more advanced medical, scientific, and consumer applications could be developed. The successful development of miniaturisation requires careful planning, and computer simulations are helping with this aspect. However, continued miniaturisation also lead to several challenges concerning this matter. It is common knowledge that the world around us is made of atoms, however for general macroscopic phenomena, the concept of a continuum works very well. Instead of looking at individual atoms, the behaviour of the collection of a very large number of atoms is studied and is expressed in mathematical equations which can be solved accordingly. On the other hand, recent decades saw a lot of interest in nanotechnology, i.e. the study of phenomenon that happen at the nano-scale (< 0.0001 mm) which means controlling of matter on an atomic and molecular scale. In these cases the concept of a continuum will fail and different techniques must be used to simulate these phenomena. However, there exist an intermediate region (where the typical size is ~ 0.00001 mm$-$ ~ 0.01 mm), where both the molecular effects and the continuum effects can be of importance. To simulate this with one single method is very challenging. For example, any method that is very accurate at the small scale will soon be too cumbersome at the large scale. On the other hand, any method that is efficient for the large scale generally lost all the details at the small scale. A solution to this problem is to use both methods at the same time and only apply it to the region where the specific method is most suitable. However, in order for this to work, these methods should communicate with each other. Effectively this means that the methods are coupled and are able to resolve the physical phenomena over a wide range of scales. The subject of this thesis is to develop, implement and test one of these methods, which are generally known as multiscale, coupled, or hybrid methods.

In the present work, the Schwarz alternating method is chosen to couple a domain that simulates a dense liquid using molecular dynamics (MD) and a domain that uses continuum methods. The method couples the two domain on the principle that the two domain solve for the same solution in a region where they overlap. Inside this overlap-region the two domain interchange boundary conditions obtained from each other. The boundary conditions for the continuum domain can easily be obtained from the MD simulation results. The specification of boundary conditions on the MD domain is less straightforward. In the current implementation it involves three main steps, which are dealt with in Section 3.2, both for the coupling of one-atom liquids like argon and more complex liquids like water. However, before the coupling can accomplished, it is necessary to obtain more information about where and when a coupling is possible and under which assumptions these coupled simulations give a correct solution. Therefore, the present work also demonstrated several MD simulations/studies to investigate the possibilities and limitations of pure MD simulations, while also the possibilities and limitations of pure continuum methods are studied.

The results of the MD simulations showed that large deviations between continuum mechanics and MD are especially noticeable near the solid walls of (nano-sized) channels or near obstacles and are local. These deviations, visible as large variations in the sampled

(continuum) macroscopic variables in the MD simulation, are the result of the interaction of the atoms in the liquid with the atoms in the solid wall and can also be observed experimentally. Only far away from the wall the macroscopic variables show their (expected) continuum value without variations. However, although these large variations indicate large non-continuum effects, even for small nano channels of about 5 nm, near continuum-like behaviour can be extracted from the results. This fact was used to determine the viscosity as a function of temperature for four different water models. In general, the results showed that in a channel with a height of about 8 nm yield very good overall continuum-like behaviour. However, this does not yet mean that a pure continuum method to compute the flow inside these channels is advisable. There are also different reasons why not to use a continuum method to simulation certain phenomena, because MD simulations do have some unique benefits, like predicting realistic wall-fluid interactions. Furthermore, with MD several other phenomena can be simulated that are difficult or even impossible with a continuum technique, like the nano-jet and nano-jet breakup. However, care must be taken, because frequently used values of the cutoff radius are too low to accurately model several important phenomena. On the other hand, comparison of continuum results with experimental result showed that, especially the continuum techniques describing electrokinetic effects, are reasonably accurate enough, even for a nano-sized device where the height is only 150 nm.

The results of several coupled simulations are shown in Section 3.3. Here it is explained how the coupled simulation can be seen as a new boundary condition for the continuum, where the value is now more accurately supplied by the communication of the MD and continuum domain. This was demonstrated by simulating Poiseuille flow of argon and water inside a large channel. The non-continuum effects near the wall are simulated accurately by MD and no expensive MD computation time is wasted on the part that resembles a continuum. The coupling of MD and continuum also enabled the specification of non-periodic boundary conditions for MD systems, which are difficult or impossible to implement in a pure MD case. This was demonstrated by a two-dimensional coupled simulation of a nanowire inside an uniform flow of argon. The main benefits and results of this type of simulation are that it investigates the influence on the flow of one obstacle, instead of the one and all its periodic images. The principles behind the coupling of domains can also be applied to other macroscopic variables than velocity, for example temperature. In this work a qualitatively study was performed on a particle inside a temperature gradient field by coupling the MD domain and the continuum domain, effectively investigating thermophoresis in liquids.

Finally, in Chapter 4 a different kind of coupling between molecules and the continuum is explained, which is very efficient to study the behaviour of polymers. For this purpose, a mesoscale simulation to measure the strength of the velocity flux needed to push a polymer into a narrow channel is demonstrated. Here excellent agreement is found with a prediction based on a de Gennes blob model of the polymer, that the critical velocity flux for translocation depends linearly on the temperature, but is independent of the length of the polymer chain or the width of the channel.

Contents

Chapter 1

Introduction

1.1 Motivation and Background

During the last two decades there is a lot of interest in numerical methods that are able to resolve physical phenomena occurring both at the very small scale as well as the larger scale. However, to accomplish this with one single method is very challenging. For example, any method that is very accurate at the small scale will soon be too cumbersome at the large scale. On the other hand, any method that is efficient for the large scale generally lost all the details at the small scale. A solution to this problem is to use both methods at the same time and only apply it to the region where the specific method is most suitable. However, in order for this to work, these methods should communicate with each other. Effectively this means that the methods are coupled and are able to resolve the physical phenomena over a wide range of scales. The subject of this thesis is to develop, implement and test one of these methods, which are generally known as multiscale, coupled, or hybrid methods.

The choice of the numerical method used for the small scale or the large scale depends on the substance and physical phenomena under interest and the way how these are coupled is linked to this. Generally speaking, the different coupled methods derived in literature can be divided into coupled simulation involving either a solid, a liquid, or a gas. In these cases, continuum mechanics is used for the large scale description, while particle or discrete methods are used to describe the small scale. For example, previous proposals for multiscale methods for solids concentrated on crack propagation in silicon [3, 164]. Here, at the large scale, the commonly employed numerical method for solid mechanics, finite element methods, are used. However, in the region near the crack the bonding between the atoms is very accurately modelled with molecular dynamics and a semi-empirical quantum-mechanical model. Other research concentrated on the modelling of crystals with more than a single grain [170]. Here, again the large scale is computed with finite element methods, while the atomic structure is modelled with a specialised (discrete) model for crystals. Once these models are coupled, they are able to very accurately predict the propagation of the crack or the behaviour of the crystal while the large scale effect are also computed. The multiscale methods for solids have the characteristic that the motion of the atoms is local, i.e. atoms do vibrate but will more or less remain in the same position.

On the other side of the spectrum there are the multiscale methods for gases. Now the atoms are generally not at the same position for very long and only occasionally interact with each other. When dealing with gases, the Knudsen number (Kn), which is the ratio between the mean free path of the gas molecule and a characteristic length (e.g. the height of a channel), gives a reasonable indication when the continuum assumption fails. If the Knudsen number is much larger than unity, the gas is in the free-molecule regime and continuum certainly fails and particle methods need to be applied. On the other hand, when the Knudsen number is much smaller than unity, the substance behaves like a viscous fluid and continuum is applicable. However, in the intermediate regime the continuum concept can still be applied, as long as the continuum equations are adapted to account for several phenomena , e.g. the now incorrect assumption of a no-slip boundary condition

at the wall. However, there are cases where this is not possible and in this situation more accurate solutions could be obtained if the particle methods and the continuum methods were to be coupled. For example, previous research on coupled methods for gases concentrates on hypersonic (rarefied) flows [26, 186, 191, 203] round a body or vehicle at high altitude. Here the Boltzmann equation, solved using a particle method like Direct Simulation Monte Carlo Method (DSMC), is the model used to accurately model the boundary layer near the body and a Navier-Stokes flow solver is used to compute the (global) flow. The two equations can be coupled by either the concept of friction [26], half-fluxes [186], or automatic domain decomposition [191]. Other types of multiscale methods for gases used an overlapped Schwarz alternating method to couple a DSMC method and a Stokes equation. This research concentrated on the simulation of a microfluidic filters for particle trapping and sorting [6], or an adaptive mesh and algorithm refinement method to couple DSMC with a compressible Navier-Stokes solver [69, 209].

Next the multiscale methods that can be used for liquids are discussed. In the case of gases there is only occasional interaction between atoms, which are free to move. In the case of solids there is very strong interaction between the atoms, which remain more or less at the same position. Between these cases, there are the atoms that represent a liquid. The characteristic of a liquid is the fact that the atoms are constantly interacting with each other while the atoms can diffuse. This also means that the type of interaction changes all the time and the Knudsen number cannot be used. Therefore, for liquids it is less obvious to tell when continuum fails. A very accurate but also computational expensive method to simulate the interatomic interaction between atoms is Molecular Dynamics (MD). Although computer power has increased dramatically recently, which means MD can be employed successfully to study important phenomena at the nano scale, like wall effects, the method is too expensive to be applied to complete nano devices. On the other hand, from an industrial point of view, the devices that are used, for example for medical applications, become smaller and smaller. In order to successfully design these devices, modelling is essential. Because the size of the device is very small, the conventional continuum techniques can fail. Also, especially in medical applications but also in more general micro and nano-fluidics, the most important physics or phenomena happen at the molecular level. But, it is impossible (and senseless) to model this device completely with a technique like molecular dynamics. Therefore, combining the two methods: molecular dynamics and conventional continuum methods, will solve this problem.

A concise review of the development of multiscale methods applied to liquids is given by Wijesinghe et al. [208]. The main conclusion of this paper is that it does not matter which solution coupling approach is used, as long as the continuum boundary condition is correctly imposed on the molecular domain. An early, but successful example, of a way to impose the continuum boundary condition was given by O'Connell et al. [40]. They introduced a (one-dimensional) method where the dynamics of the atoms, near the boundary where MD and the continuum are coupled, are constrained. The coupled method was demonstrated to work by simulating unidirectional startup flow of a simple fluid near a solid interface. However, the disadvantage of the solution coupling approach they used, is the fact that the timescale of the MD domain and continuum domain is coupled, which is a major limitation. Hadjiconstantinou et al. [81, 82] therefore proposed a solution coupling approach based on a domain decomposition technique known as the Schwarz alternating method. By using this method, one domain is subdivided into several subdomains. Length scale decoupling is achieved by using an overlap area between each subdomain, while timescale decoupling is achieved automatically. The reason for this is the fact that the Schwarz technique is a steady-state solution method. However it is still possible to simulate transient problems [81, 82]. Note that the use of an overlap area does mean that both methods that are coupled should be able to compute the same solution inside the overlap area. Their method worked on matching the velocity inside the overlap area and was validated by simulating a moving contact-line problem. Werder et al. [206] also used the Schwarz alternating method to couple the velocity field of a domain using molecular dynamics and a domain using a continuum method, but expanded the coupled framework to a two-dimensional one. In their work they accentuate the fact that large disturbances to the fluid structure are visible near the overlap region, if nothing is done to prevent so. These disturbances can be minimised by introducing an effective boundary potential that mimics the missing interactions near the overlap domain. In more recent years, Kotsalis et al. [110] extended this boundary potential. Another way to couple the MD and continuum domain using a domain

decomposition technique, is to couple fluxes. For example, Flekkoy et al. [61] introduced a method that aimed at keeping the fluxes of the conserved quantities continuous across the MD-continuum interface. Nie et al. [137] introduced an improved version and showed how it can be applied to an impulsively started Couette flow. However, by using the direct flux exchange method, the time scales of the molecular and continuum simulations are again coupled [82, 160], which, as mentioned before, is a major limitation. Further, Werder et al. [206] showed that in order to obtain statistically meaningful estimates of the fluxes for the continuum domain, an excessive amount of sampling of the atomistic region is needed, which further adds to the computational effort needed for the flux method. The coupled methods discussed so far were based on the fact the molecular scale and macroscopic scale were computed in separate domains and coupled by means of an overlap area, i.e. using domain decomposition methods. However, there is another strategy described in literature, i.e. the so-called heterogeneous multiscale method (HMM) [53]. The idea of HMM is to use a macroscopic model and numerical mesh for the whole domain, i.e. the part where continuum is valid but also the part where a continuum method fails. However, it is the task of a separate microscale model to solve for the non-continuum part locally and provide meaningful macroscale data that is used/interpolated on the macroscopic model. A detailed review of the heterogeneous multiscale method is given by E et al. [54], while a comparison between this method and the other methods is given by Ren [159].

In the present work, the Schwarz alternating method is chosen to couple a domain that simulates a dense liquid using molecular dynamics and a domain that uses continuum methods. The main reasons the Schwarz alternating method is chosen are:

- Length scale decoupling is achieved by using an overlap area between each subdomain, while timescale decoupling is achieved automatically
- Compared to other MD/Continuum coupling methods, this method requires less amount of sampling of the atomistic region in order to obtain statistically meaningful estimates of continuum values
- The fact that two separate domains can be solved and only the overlap region between the two domains need to be matched is an advantage.

The reason why this last item is an advantage is the fact that it will enable the possibility to not only couple the macroscopic variable of velocity, but also other variables like temperature or stress, either separately or combined. Further, previous research mainly demonstrated how coupling can be achieved and illustrated the method by doing simulations where the liquid consisted of simple (noble) atoms, i.e. no molecules. Here, it also shown how a simulation with molecules, especially water, can also be used in a coupled approach. On the other hand, recent developments in MD and computer hardware means more can be achieved using MD alone. Very interesting phenomena that occur at the nanoscale, can be modelled with MD very accurately. However, the effects on the larger scale in these simulations must be ignored because of the computational limitations. A successful coupling of MD and a continuum method might resolve this and a better understanding of these phenomena is the result. However, before this can be accomplished, it is necessary to obtain more information about where and when a coupling is possible and under which assumptions these coupled simulations give a correct solution. Therefore, the present work also demonstrates several MD simulations/studies to investigate the possibilities and limitations of pure MD simulations, while also the possibilities and limitations of pure continuum methods in similar conditions are studied. The results of these simulations are essential for coupling and verification purposes. The Schwarz alternating method only works if the overlap area between the two domains solve the same solution. Therefore a pure MD simulation will reveal where the continuum method could also be used, while it also provides a verification for the coupled solutions.

1.2 Outline and Scope of the Thesis

The focus of this thesis is the development of a coupled numerical simulation approach for dense liquids where one domain is represented with molecular dynamics (MD), while the other other domain solves the conventional continuum equations. The first part of the thesis will focus on how well the continuum methods can be applied to nano-sized

devices, followed by how MD can be used to provide better answers. Also, the limits of MD simulations are explored, especially how a commonly used trick in the MD code to decrease the computational effort, can greatly affect the final solution. The MD simulations are also performed to investigate when MD is really necessary or when conventional continuum techniques would provide the same answer. The second part will focus on how the two methods can be coupled and how unique (better) solutions can be obtained if both methods are used in a coupled manner.

The main contents of this thesis is divided in three chapters. The first one, Chapter 2, has the purpose to introduce the reader to the numerical models that are used throughout the book, but at the same time focuses on the phenomena encountered in nano-sized devices. In Section 2.1 these phenomena, like electrokinetic flows and thermophoretic flows, are introduced. The section is followed by Section 2.2.2 where the numerical technique to solve the continuum equations is explained, namely the Finite Element Method (FEM). On the other hand, Section 2.3 is entirely devoted to the numerical technique to solve the interatomic interactions, namely Molecular Dynamics (MD). Chapter 2 ends with several applications and results obtained either with a pure continuum method or completely with MD. Here it is addressed that for relatively large nano-sized devices ($\sim$150 nm) continuum techniques can still be employed. However, for smaller devices ($<$60 nm) or for devices where the most important feature is on the molecular scale, e.g. the influence of one DNA molecule or wall effects, these pure continuum techniques can fail. The remainder of Section 2.4 focuses on employing only MD to investigate these phenomena.

Chapter 3 has the purpose to introduce a method how the two domains, continuum and molecular, can be coupled. In Section 3.1 the methodology of the Schwarz alternating method is introduced, while Section 3.2 goes in more detail how this technique is employed. The chapter ends with Section 3.3, where several applications and results are shown of coupled simulations.

Chapter 4 shows how a different type of coupling between the molecular and the continuum domain can be used to investigate the flow injection of polymers into nanopores.

Finally, in Chapter 5, the main conclusions are summarised and recommendations for further work and applications are given.

Chapter 2

Numerical Models

2.1 Modelling the Continuum: Computational Fluid Dynamics (CFD)

Throughout this thesis the terms "molecular dynamics", "continuum", and "non-continuum" are used. In section 2.3 it will be explained what is meant with molecular dynamics, but first the term continuum is further specified. When something is called a continuum it is assumed that the substance of something is distributed throughout and has no gaps or holes, which means that it can be divided infinitely. A continuum is most of the times simply referred to as a solid, liquid, or a gas. This also means that the continuum completely ignores the fact that everything is made of molecules. On the (classical) molecular scale, a solid, a liquid, or a gas, can consist of the *same molecules* or *atoms* but with different spatial configuration, where the rest of the space is a vacuum. The good thing about the concept of a continuum is that it enables the possibility to express the behaviour of the substance in differential equations that, when correctly applied, give the same results as experiments. The term "non-continuum" basically refers to the situation where all or some of the above is not applicable. However, to tell *when* this is the case is tricky, which is discussed next.

When dealing with gases, the Knudsen number (Kn), which is the ratio between the mean free path of the gas molecule and a characteristic length (e.g. the height of a channel), gives a reasonable indication when the continuum assumption fails. If the Knudsen number is much larger than unity, the gas is in the free-molecule regime and continuum certainly fails. On the other hand, when the Knudsen number is much smaller than unity, the substance behaves like a viscous fluid and continuum is applicable. However, in the intermediate regime the continuum concept can still be applied, as long as the continuum equation take several phenomenon like the now incorrect assumption of the no-slip wall into account. Of course continuum still fails when the molecular effects are becoming important in this intermediate regime.

For liquids it is less obvious when continuum fails. The Knudsen number cannot be used, because in the liquid state the molecules are in close contact and constantly interacting with each other. Therefore for a liquid the question can be: how many molecules are needed before the statistical fluctuations are small enough to accurately specify a mean value and what typical length is associated with that? The statistical fluctuations d are related to the number of atoms N as $d = 1/\sqrt{N}$ [103], while the density of the liquid controls the occupied volume of the N atoms. For example, for water at standard conditions, when the statistical fluctuations are less than 1%, the typical length is 6.7 nanometres. However, this length is only an indication again, because on this small scale other effects than atomic interaction like electrical effects can be very strong. Also the presence of a solid structure such as a channel wall can influence the continuum features of the liquid. In such cases molecular dynamics is very effective to use. In section 2.4.2 an estimate is given for the minimum size of a (nano-sized) channel where continuum techniques can safely be applied, using molecular dynamics.

From the above description of a continuum it is clear that as long as the typical dimension of the flow is much larger than any molecular effect, the continuum approach can safely be applied. More general, this can be specified by a surface to volume ratio, indicating the importance of surface effect compared to volume effects. If this ratio is high, the surface effects dominate and the continuum approach is very likely to fail. In the case when microfluidics are considered, the typical dimension is much larger than any molecular effect. Here the definition of microfluidics is: the description of the flow and phenomena of fluids, where at least one of the dimensions is in the micrometre scale, usually defined as < 100 micrometre. In microfluidics the flow can safely be described with the Navier-Stokes equation, however other phenomena, for example surface effects like electrokinetics, also become important because of the small scale. For the modelling of the flow this means that these effects must be included in the simulation. Section 2.1.2 introduces the most frequently observed and used phenomena, which are electroosmosis, electrophoresis, and di-electrophoresis. Beside the electric effects, another phenomenon can be observed when temperature gradients are present. Again, because of the small scale, small temperature differences inside a channel result in very large temperature gradients inside the channel. These temperature gradients are the source of another phenomena called thermophoresis, which will be explained in section 2.1.3. In general, these phenomena become more important when the typical dimension becomes smaller i.e. when the surface to volume ratio becomes higher. The regime where the typical dimension goes down to nano size is called nanofluidics. In this case the continuum hypothesis is not necessarily valid and surface effects (e.g. confinement in the form of a channel) are now dominating. This also means that if you want to study small scale flow and include molecular effects (e.g. solid boundaries) you have to use a coupled method where molecular dynamics and continuum is solved together. How the coupling can be achieved, is discussed in chapter 3. Luckily in most cases the normally used continuum equations are still applicable with some minor alterations which will be discussed in the next subsections. This section ends with a overview of the continuum numerical methods used throughout the book.

2.1.1 Continuum Flows (Navier-Stokes)

In the introduction it was described what continuum entails and that it can be used to successfully describe microfluidics and nanofluidics. In this section the basic continuum (governing) equation for fluid flow are given. In typical microfluidic and nanofluidic applications, the flow can be considered as incompressible. In this case the conservation of mass can be specified as:

$$\nabla \cdot \vec{u} = 0 \tag{2.1}$$

where $\vec{u}$ is the velocity vector. Equation 2.1 is called the *continuity equation*. The second equation is the conservation of momentum, which equals Newton's second law applied to a small volume element of fluid, and can be written as:

$$\rho_f \left(\frac{\partial \vec{u}}{\partial t} + (\vec{u} \cdot \nabla) \vec{u} \right) = \vec{f_b} + \vec{f_s} \tag{2.2}$$

where ρ_f is the density of the fluid, $\vec{f_b}$ is a *body force density*, and $\vec{f_s}$ is a *surface force density*. The surface force density for a Newtonian fluid with constant viscosity μ_f is given by:

$$\vec{f_s} = -\nabla p + \mu_f \nabla^2 \vec{u} \tag{2.3}$$

where p is the pressure. The i-th component of the body force density, $f_{b,i}$, in terms of a general stress tensor σ_{ij} is given by:

$$f_{b,i} = \frac{\partial \sigma_{ij}}{\partial x_j} \tag{2.4}$$

which in the case of gravity is simply given by $\vec{f_g} = \rho_f \vec{g}$. Substitution of all the equations into equation 2.2 results in:

$$\rho_f \left(\frac{\partial \vec{u}}{\partial t} + (\vec{u} \cdot \nabla) \vec{u} \right) = -\nabla p + \mu_f \nabla^2 \vec{u} + \rho_f \vec{g} \tag{2.5}$$

which is called the *Navier-Stokes equation* and is valid for any isotropic, incompressible flow with constant viscosity. However, in many cases discussed hereafter, there is no need to solve the full Navier-Stokes equation. The flow is these cases are considered steady and laminar and therefore can safely be described with the *Stokes equation*:

$$\nabla p = \mu_f \nabla^2 \vec{u} + \rho_f \vec{g} \tag{2.6}$$

2.1.2 Electrokinetic Flows

In this section the electrokinetic effects and the resulting electrokinetic flow are discussed. These electrokinetic effects basically happen on the molecular scale, but luckily they can be reasonably described with pure continuum equations only.

When a micro or nano channel is filled with a liquid, especially a polar liquid like water, the walls of the channel acquire a surface electric charge. This charge is caused by natural processes like ionisation, ion adsorption, or ion dissolution. However, in practical applications the liquid inside the channel also contains a concentration of certain salts, with an equal amount of negatively and positively charged ions. Because the wall also has a certain charge, the ions of opposite wall charge (counter-ions) are attracted towards the surface, whereas ions of like charge (co-ions) are repelled away from the surface. However, due to thermal motion, the ions tend to distribute themselves evenly over the solution. The result is a compromise in which the concentration of the counter-ions is high near the surface and decreases gradually till the bulk value is reached at a large distance from the surface. On the other hand, close to the surface a rigid charge layer of counter-ions is formed. This layer, which has the thickness of approximately one ion, is called the Stern layer. This layer is bounded by strong electrostatic forces, and hence it is immobile. The Stern layer together with the diffuse layer of co-ions and counter-ions, is called the electric double layer (EDL). The typical thickness of this layer is determined by the opposing forces of electrostatic attraction and thermal diffusion, which is on the order of the Debye length of the medium. Typical values range from one nanometre to several hundreds of nanometres and therefore can ignored in macroscopic flows, but is important in microfluidics and nanofluidics, especially when the typical scale becomes comparable to the Debye length.

So far the theory still involves discrete objects like ions and molecules and the basic principles all happen at this molecular scale. However, the fact that there is an electric double layer results in some very useful effects on the much larger (continuum) scale. For example, when an external electrical field is applied to the electric double layer, this will set the ions inside the diffuse layer in motion. This motion will induce drag due to the fluid viscosity, which as a result induces motion of the (continuum) bulk fluid. This principle is called *electroosmosis*, and the resulting flow is called *electroosmotic flow* (EOF). This flow can be used to transport fluids inside a micro or nano sized fabricated device. However, this same electric double layer is also formed at the solid-liquid interface of a solid (charged) particle when this particle is immersed in the liquid present in the device. When in that case an external electric field is applied, the charged particle and the ions inside the diffuse layer are set in motion. This principle is called *electrophoresis* and can be useful to transport individual particles. Another useful method for the manipulation of particles is called dielectrophoresis, which is the motion of particles caused by polarization effects in an applied *non-uniform* electric field. In this case the particle does not have to be charged, and it only occurs whenever there are gradients inside the electric field, which can be either AC or DC. Dielectrophoresis is especially useful for the manipulation of molecules or (biological) cells.

The EDL can be utilised for continuum applications, therefore to predict the continuum behaviour, the governing equations need to be specified. These are discussed in detail in appendix A but the general idea is as follows: For electroosmotic flow, the governing equations are four coupled equations [123]. The first equation specifies the transport of the ions in the solution, which is the Nernst-Planck equation. The second equation, derived from Maxwell's equations, specifies the relation between the electric field and the net charge density of the ions obtained from the concentration of the ions. The third and fourth equations specify conservation of mass and momentum, which are the continuity equation and the Navier-Stokes equation discussed before. However, the latter is modified such that the interaction between the ions and the applied electrical field are taken into account. Fortunately, in most cases there is no need to solve all four equations simultaneously, and simplifying approaches can be used, which are summarised in section 2.1.2.1. The same

applies for electrophoresis and dielectrophoresis, where in most cases the governing equations can be reduced to a simple force balance between the (di)electrophoretic force and the Stokes drag force, which are given in section 2.1.2.2 and 2.1.2.3. respectively.

2.1.2.1 Electroosmosis

The first electrokinetic effects that will be described is electroosmosis and electroosmotic flow, where the charged ions inside the electric double layer near the channel wall are used to set the bulk fluid into motion. The governing equation for electroosmotic flow are given in appendix A. There it is explained that the typical thickness of the EDL, given by the *Debye length*:

$$\frac{1}{\kappa} = \sqrt{\frac{\epsilon\epsilon_0 k_b T}{2z^2 e^2 c_\infty}} \tag{2.7}$$

compared to the typical dimension of a microfluidic or nanofluidic device can be used to greatly simply the equations. In the case of microfluidic devices, the typical dimension of the device, h, is on the order of several micrometres while the typical thickness of the EDL is in the nanometre scale. This case is called the thin EDL case and the EOF velocity is given by the *Helmholtz-Smoluchowski velocity*:

$$\vec{u}_{EOF} = -\frac{\epsilon\epsilon_0\zeta_w}{\mu}\vec{E} = \alpha_{os}\vec{E} \tag{2.8}$$

where $\alpha_{os} = -\epsilon\epsilon_0\zeta_w/\mu$ is a constant that depends only on the surface and electrolyte properties. The Helmholtz-Smoluchowski velocity is the velocity at the outer edge of the EDL, induced by the moving ions inside the EDL. Because the EDL is very thin, from a macroscopic point of view this can be seen as a slip velocity specified at the wall. This also means that the electric potential is effectively only the externally applied potential ϕ and which can be solved using a Laplace equation:

$$\nabla^2\phi = 0 \tag{2.9}$$

where the boundary conditions can be an applied voltage by means of a Dirichlet boundary condition or an electric insulation/leakage boundary condition by means of a Neumann boundary condition. The electric field is derived by taking the minus gradient of the potential: $\vec{E} = -\nabla\phi$. However, under certain conditions [42] the process of specifying the slip velocity and solving for the flow can be skipped altogether, in which case the local velocity in the bulk is entirely given by the Helmholtz-Smoluchowski relation. In that case the flow pathlines exactly follow the electric field lines.

On the other hand, in a nanofluidic device the dimensions of the device are similar to the size of the EDL or smaller. In this case the full theory still needs to be applied. However, when the typical thickness of the EDL is several hundreds of nanometres, the EDL can be considered "thick" compared to the nanodevice. In this case, i.e. the thick EDL case, the EOF velocity can be derived using a Hele-Shaw-like approximation [86]. Consequently, the velocity profile is a Poiseuille velocity profile and the average EOF velocity is given by the expression:

$$\langle\vec{u}_{EOF}\rangle = -\frac{\epsilon\epsilon_0\zeta_w}{\mu}\vec{E}\frac{(\kappa h)^2}{12} \tag{2.10}$$

which is just the Helmholtz-Smoluchowski relation multiplied with a certain factor given by the value of κh.

2.1.2.2 Electrophoresis

In the last section it was shown how the EDL that is formed near the walls of micro or nano sized fabricated device can be used to create electroosmotic flow inside the device. However, when a solid (charged) particle or bead is inserted inside this device, an electric double layer is also formed around the particle. When an external electric field is applied, the charged particle and the counter-ions inside its diffuse layer are set in motion. This principle is called electrophoresis [123]. Similar to what was the case with electroosmotic flow, when a very thin or a very thick electric double layer compared to the particle radius is considered, simple expression for the electrophoretic velocity can be derived[123]. In the

case of a thin EDL the derivation is similar to that of the Helmholtz-Smoluchowski relation, and the value of the electrophoretic velocity can be determined with the same equation except for the sign. In the case of a thick EDL the particle can be seen as point charge in an electric field. The electrophoretic velocity is then given by the Hückel equation, which is 2/3 the value of the velocity in the thin EDL case. In the general case, the fact that the charged particle and the counter-ions inside the diffuse layer have an opposite charge and thus opposite direction of motion, must be taken into account. Eventually the velocity of the particle will be somewhere between the value of the thin and thick EDL case. Therefore the general electrophoretic velocity is given by [123]:

$$\vec{u}_{EP} = \frac{2}{3} f\left(\kappa r_p\right) \frac{\epsilon \epsilon_0 \zeta_p}{\mu} \vec{E} \tag{2.11}$$

where $f\left(\kappa r_p\right)$ is called the *Henry function* which is a dimensionless function that is dependent on the ratio of particle radius, r_p, and to double layer thickness. This function goes from 1.0 in the thick-EDL case to 3/2 in the thin-EDL case.

2.1.2.3 Dielectrophoresis

So far two electrokinetic principles were discussed which were based on the principle that the wall of the device and the particle were charged. However, there is another very interesting and useful electrokinetic principle that does not require the particle to be charged at all and that is called di-electrophoresis [151]. This principle is based on the fact that when a particle is polarised due to an electric field, the particle forms a dipole. When the electric field would be uniform, nothing would happen, because the force on each pole of the di-pole is the same and opposite in direction. However, when this same particle is put in a *non-uniform electric field*, at each pole a different force is applied, and hence a net force acting on the particle remains that can set the particle into motion. In general the di-electrophoretic force on the particle can be expressed as:

$$\vec{F}_{DEP} = \vec{p} \cdot \vec{\nabla} \vec{E} \tag{2.12}$$

where $\vec{p}$ is the dipole moment vector and $\vec{\nabla}\vec{E}$ is the gradient of the electric field. For the special case of a homogeneous dielectric sphere with radius r_p and permittivity ϵ_2 immersed in a fluid with permittivity ϵ_1, the expression for the effective dipole moment is [100]:

$$\vec{p}_{eff} = 4\pi\epsilon_1 K\left(\epsilon_2, \epsilon_1\right) r_p^3 \vec{E} \tag{2.13}$$

where:

$$K\left(\epsilon_2, \epsilon_1\right) = \frac{\left(\epsilon_2 - \epsilon_1\right)}{\left(\epsilon_2 + 2\epsilon_1\right)} \tag{2.14}$$

is called the Clausius-Mossotti function that provides a measure of the strength of the polarisation of the spherical particle as a function of the permittivities of the particle and the fluid. Substitution of equation 2.14 into equation 2.13 and making use of a simple vector identity, the di-electrophoretic force for the sphere equals:

$$\vec{F}_{DEP} = 2\pi\epsilon_1 K\left(\epsilon_2, \epsilon_1\right) r_p^3 \vec{\nabla} E^2 \tag{2.15}$$

The di-electrophoretic velocity of the sphere can be determined by the fact that when the motion of the spherical particle is steady, the di-electrophoretic force is equal to the Stokes drag force, hence:

$$\vec{u}_{DEP} = \frac{\epsilon_1 K\left(\epsilon_2, \epsilon_1\right) r_p^2}{3\mu_f} \vec{\nabla} E^2 \tag{2.16}$$

where E^2 is the electric field intensity. Traditionally, two di-electrophoretic classes are distinguished; positive di-electrophoresis (+DEP) and negative di-electrophoresis (-DEP). In the case of +DEP, $K > 0$ or $\epsilon_2 > \epsilon_1$, and the particles are attracted to electric field intensity maxima and repelled from minima. On the other hand, for -DEP, $K < 0$ or $\epsilon_2 < \epsilon_1$, and the particles are attracted to electric field intensity minima and repelled from maxima.

This concludes the section where three electrokinetic effects were discussed that can be used for useful applications in microfluidics and nanofluidics; electroosmosis, electrophoresis, and di-electrophoresis. In section 2.4 an example will be given of a nano-sized device where all three effects are used to move and sort DNA molecules.

2.1.3 Thermophoretic Flows (Thermophoresis)

In the last section several electrokinetic effects were discussed that can be used to move or manipulate particles or DNA molecules inside microscopic or nanoscopic sized devices using electric fields and the response of the particles inside this electric field. However, in the last two decades another effect has attracted more attention for this purpose too. The driving force behind this effect is the temperature (gradient) field which results in a mass flux and effectively particles or species will migrate. This effect is called thermophoresis, or alternatively known as thermal diffusion, thermodiffusion, or the Ludwig–Soret effect [105] after Ludwig who discovered the effect in 1856 and Soret who further investigated it in 1897. The purpose of this section is to give a simple introduction to thermophoresis. For more details see for example [105, 150, 207] and the references therein.

Similar to the derivation of the electrokinetic effects, it starts with the Nernst-Planck equation, equation A.3 which sums several types of transport of species in a (dilute) solution. Thermophoresis is simply an additional particle transport type inside the equation for the total mass flux. In the case where there is no electric field or external flow, and low concentration, this mass flux is written as [150]:

$$J = -D\nabla c - cD_T\nabla T \tag{2.17}$$

where c the concentration of the species in solution, D denotes the diffusion coefficient, and D_T is called the thermal diffusion coefficient (or better: thermophoretic mobility). The first step in investigating thermophoresis further is to consider the stationary state where the mass flux J disappears. Equation 2.17 can be rewritten to:

$$S_T = -\frac{1}{c}\frac{\nabla c}{\nabla T} \tag{2.18}$$

where $S_T = {}^{D_T}/_{D}$ is called the *Soret coefficient* and gives an indication how effective the thermophoretic effect is. On the other hand, a steady-state concentration gradient is obtained, described by:

$$\nabla c = -cS_T\nabla T \tag{2.19}$$

which means that depending on the sign of the Soret coefficient, particles focus either at the cold or the hot side. This is also indicated by the the drift velocity of a particle inside the temperature gradient, given by :

$$u_T = -D_T\nabla T \tag{2.20}$$

which means that when $S_T > 0$ the particles move to the cold side, while when $S_T < 0$ the particles move to the hot side. It is however the sign and value of the Soret coefficient itself which is difficult to predict, as for example, it can change sign, simply by changing the concentration, temperature, pH, or ionic strength of the solution. For gases there exist a theory that describes thermophoresis very well [78], however for liquids this generally is not the case and the value and correct sign of the Soret coefficient are not predicted. Molecular dynamics simulations proved to be useful for the investigation of thermophoresis and showed both good qualitative and quantitative agreement with experiments [207]. However, several continuum approaches are applicable to larger molecules or colloidal suspensions.

In the case of colloidal suspensions, applied to thermophoresis, a similar method how electrophoresis was explained in section 2.1.2.2 can be used. In the case of electrophoresis, an electric double layer is formed around any particle, and when an external electric field (a potential *gradient*) is applied, the counter-ions inside its electric double layer are set in motion. In other words, the resulting *bulk* motion is the result of movement of ions in a very small layer near the particle only. It was also noted that the resulting electrophoretic velocity is generally a function of several properties, summarised in the Henry function, $f(\kappa r_p)$. In the limit of a very small or very large diffuse layer compared to the particle radius, the velocity becomes independent of this function. Connected to this, Ruckenstein [163] showed how thermophoresis, diffusiophoresis and electrophoresis can be described as the result of a variation of the surface tension. Starting with the electrophoretic velocity in the case of a thin EDL and one dimensional case, equation 2.11 can be rewritten to:

$$u_{EP} = -\left[\frac{\epsilon\epsilon_0\zeta_p}{\sigma}\right]\frac{1}{\mu}\frac{d(\sigma\phi)}{dz} \tag{2.21}$$

where σ is the surface charge density. The next step is to notice that the term between square brackets has the dimension of a length, while the last factor is the gradient of a kind

of interfacial tension due to the external field. Therefore the equation can be rewritten to:

$$u_{EP} = -\frac{l}{\mu}\frac{d\gamma}{dz} \tag{2.22}$$

where under the conditions of a low surface potential, the characteristic length is determined by the reciprocal Debye length, $l = \kappa^{-1}$ and γ is the interfacial tension. Rubenstein suggests that a similar expression for the velocity can be written for thermophoresis, where the gradient now is due to the temperature gradient:

$$u_T = -\frac{l_T}{\mu}\frac{d\gamma}{dT}\frac{dT}{dz} \tag{2.23}$$

where $l_T = (2\kappa)^{-1}$. In other words, the motion is therefore seen as a Marangoni effect, where the interfacial tension gradients in the small layer near the particle are caused by an external field, which in the case of thermophoresis are the temperature differences.

This concludes this section where thermophoresis was shortly introduced. Although this phenomenon is sharing some of the properties of electrokinetic flows, i.e. the phenomenon is local and of molecular nature while the effect is noticeable in the bulk, there are no complete satisfactorily theories available for liquids like there are for electrokinetic flows. One of the methods that can be used to derive these models can be molecular dynamics.

2.2 Numerical: Finite Elements

The previous sections introduced the governing continuum equations of several types of flow that are commonly used or encountered when dealing with micrometre or nanometre sized devices. Roughly speaking, the equations solve for the velocity, pressure, electric, and the temperature fields, and may or may not be coupled. Of course these continuum equations need to be numerically solved, and the choice of method here is the finite element method (FEM). Although fluid dynamics tend to be mainly solved with finite volume methods, because of the automatic conservation properties, there are also good techniques available for FEM. Furthermore, there is also the need to solve for the other variables which can be done relatively easy in FEM.

The history of the finite element method mainly starts in the study of solid mechanics (e.g. structural analysis problems) as early as the 1940's. The basic idea is to discretise the continuum problem into a mesh made out of discrete sub-domains, so-called elements. The most commonly used finite element method is based on the theory of Galerkin's weighted residual. A short introduction to the theory of this method is given in the next section. For a more thorough introduction, the reader is referred to for example [52] or [217]. After the basic theory, it will be shown how the finite element method can be used to compute the temperature field and electric fields. The computation of the flow (velocity field and pressure) needs some more attention and will be discussed last.

2.2.1 Introduction / Theory

The (Galerkin) finite element method basically consist of two separate steps. First the boundary value problem or differential equation is rewritten (from the strong form) to the weak form, or variational form, which in most cases can be done by hand instead of computer. The second step involves the discretisation, where the weak form is subdivided into elements. However, to define the weak form of the boundary value problem, two collections of functions need to be defined first; the *test* or *weighting functions* and the *trial* or *admissible solutions*. To explain this further it is necessary to define a general computational domain Ω where part of the boundary, Γ_D, is specified with Dirichlet boundary conditions, and where the other part of the boundary, Γ_N, is specified with Neumann boundary conditions. Both the weighting functions and trial solutions are general functions which are (at least) square integrable and have square integrable first derivatives over the computational domain Ω. However, the weighting functions, hereafter noted as w, vanish on the Dirichlet part of the boundary, Γ_D, while the trial solutions actually satisfy the boundary conditions at the same boundary. When the variable that needs to be solved is properly chosen, this last fact can be used to automatically satisfy the Dirichlet boundary condition. However, the same thing can

also be done for the Neumann boundary condition, which is best shown with an example.

For this example, consider the Poisson equation which is solved for the computational domain Ω introduced above:

$$\begin{aligned} -\nabla^2\varphi &= f && \text{in } \Omega \\ \varphi &= \varphi_D && \text{on } \Gamma_D \\ \vec{n}\cdot\nabla\varphi &= \varphi_N && \text{on } \Gamma_N \end{aligned} \tag{2.24}$$

Here f is a source term, φ_D is the value of φ on the Dirichlet boundary, while the φ_N is the value of φ on Neumann boundary, where $\vec{n}$ is the outward normal of this boundary. The formulation of the weak form of the boundary value problem can be achieved by first multiplying the differential equation by the weighting function w and then integrating over the domain Ω which results in:

$$-\int_\Omega w\nabla^2\varphi d\Omega = \int_\Omega w f d\Omega \tag{2.25}$$

The next thing to do is to expand the first term by using some vector identities and by applying the divergence theorem to transform a volume integral to a surface integral, which results in:

$$-\int_\Omega w\nabla^2\varphi d\Omega = -\int_\Omega \left(\nabla\cdot(w\nabla\varphi) - \nabla w\cdot\nabla\varphi\right) d\Omega = \int_\Omega \nabla w\cdot\nabla\varphi d\Omega - \int_\Gamma w\left(\vec{n}\cdot\nabla\varphi\right) d\Gamma \tag{2.26}$$

Please note that by doing so, the term inside the surface integral contains the Neumann boundary condition defined before. The weighting function, w, is per definition zero at the Dirichlet boundary, which results that the surface integral is only over the Neumann part of the boundary. Consequently this means that the Neumann boundary condition is included naturally. The total equation in weak form then becomes:

$$\int_\Omega \nabla w\cdot\nabla\varphi d\Omega = \int_\Omega w f d\Omega + \int_{\Gamma_N} w\varphi_N d\Gamma \tag{2.27}$$

However, this equation does not include the Dirichlet boundary condition yet, but this can be achieved by making sure that the mathematical space where φ belongs to is the same as the space of the trial solutions, as these do satisfy the boundary condition per definition. This ends the first step for making the differential equation ready for the finite element method.

The second step involves the discretisation of the weak form. As may be obvious from the name of the method, the discretisation is done in *elements*, however an element entails more. Every element is built out of *nodes*, or *nodal points* where each element can have a different number of nodes, and several elements may use the same node. This also means that the number of elements meeting at a node may vary from node to node, which means that the method handles unstructured meshes naturally. The interpolation between all the nodes is governed by so-called *shape function*s, which will be explained in more detail later. The goal of the discretisation is to first compute the (local) contribution of each element to the integral in the weak form. After this is done, all these contributions are *assembled* to create one (global) complete system in the form of a matrix equation which can be solved by any appropriate matrix solver. By using the interpolation by shape functions as described above, the approximate solution noted as φ^h can be written as:

$$\varphi^h(\vec{x}) = \sum_{A\epsilon\lambda\setminus\lambda_D} N_A(\vec{x})\,\varphi_A + \sum_{A\epsilon\lambda_D} N_A(\vec{x})\,\varphi_D(\vec{x}_A) \tag{2.28}$$

Here λ represents a global node number in the finite element mesh, λ_D represents the nodes where Dirichlet boundary conditions are specified, N_A is the shape function associated with node number A, and φ_A is the (unknown) value at node number A. Additionally, in the Galerkin formulation the (discrete) arbitrary test functions, w^h, are defined such that they are the collection of all (finite) linear combinations of the elements of N_A. Substitution of this into the weak form, equation 2.27, and rearranging results into the following discrete

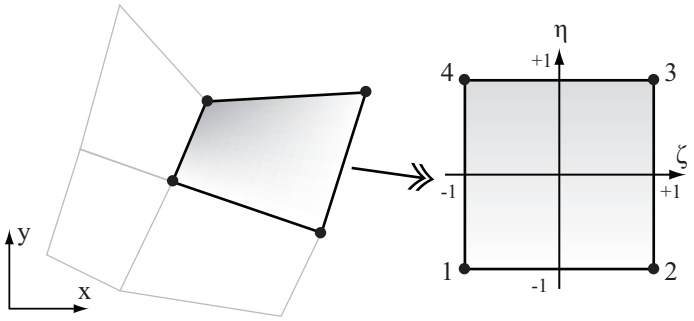

Figure 2.1: The general quadrilateral element and how the global nodes and coordinate system are mapped onto a local nodes and coordinate system.

weak form:

$$\sum_{A\epsilon\lambda\backslash\lambda_D} \varphi_A \int_\Omega \nabla N_B \cdot \nabla N_A d\Omega = \int_\Omega N_B f d\Omega + \int_{\Gamma_N} N_B \varphi_N d\Gamma - \sum_{A\epsilon\lambda_D} \varphi_D\left(\vec{x}_A\right) \int_\Omega \nabla N_B \cdot \nabla N_A d\Omega \tag{2.29}$$

where A and B represent the global node numbers and $B\epsilon\lambda\backslash\lambda_D$. What remains now is to specify the shape functions, which is done next.

The theory explained so far can be applied in 1D, 2D, or 3D, however most of the upcoming continuum simulations only ask for a 2D solution. Commonly used meshes in two dimensions consist of triangular and/or quadrilateral elements. The left illustration in figure 2.1 shows a general quadrilateral element with its *global* node numbers and attached global coordinate system. However, constructing a shape functions in the global coordinates (x, y) for this general quadrilateral would result in very complex algebraic expressions. Therefore it is better to construct the shape function in a *local* coordinate system which is shown in the right illustration. For this purpose, each quadrilateral is mapped onto a canonical square with normalised local coordinates $(\xi, \eta) \in [-1, 1] \times [-1, 1]$. The shape function of the quadrilateral in local coordinates simply is defined by:

$$N_1 = \frac{1}{4}(1-\xi)(1-\eta) \tag{2.30}$$

$$N_2 = \frac{1}{4}(1+\xi)(1-\eta) \tag{2.31}$$

$$N_3 = \frac{1}{4}(1+\xi)(1+\eta) \tag{2.32}$$

$$N_4 = \frac{1}{4}(1-\xi)(1+\eta) \tag{2.33}$$

Where the value of the variable φ can now be evaluated (interpolated) anywhere inside the mapped quadrilateral element using the nodes by:

$$\varphi^h(\xi, \eta) = \sum_{a=1}^{4} N_a(\xi, \eta)\, \varphi_a \tag{2.34}$$

Under special circumstances the same shape functions can also be used to specify the relation between the global (x, y) and local (ξ, η) coordinates, in which case the element is called isoparametric of which the four-node quadrilateral is an example. The coordinate transformation is then defined as:

$$\left\{ \begin{array}{c} x \\ y \end{array} \right\} = \sum_{a=1}^{4} N_a(\xi, \eta) \left\{ \begin{array}{c} x_a \\ y_a \end{array} \right\} \tag{2.35}$$

where a define the local nodes of the element. The discrete weak form not only uses the shape functions N, but also the derivatives with respect to the global coordinates (x, y). Luckily these derivatives can easily be converted from one coordinate system to another by

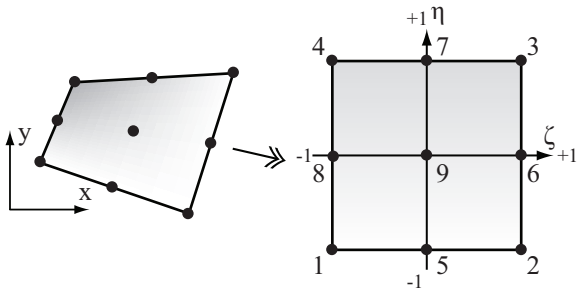

$$
\begin{aligned}
N_1 &= \tfrac{1}{4}(\xi-1)(\eta-1)\xi\eta \\
N_2 &= \tfrac{1}{4}(\xi+1)(\eta-1)\xi\eta \\
N_3 &= \tfrac{1}{4}(\xi+1)(\eta+1)\xi\eta \\
N_4 &= \tfrac{1}{4}(\xi-1)(\eta+1)\xi\eta \\
N_5 &= -\tfrac{1}{2}\left(\xi^2-1\right)\eta(\eta-1) \\
N_6 &= -\tfrac{1}{2}\xi(\xi+1)\left(\eta^2-1\right) \\
N_7 &= -\tfrac{1}{2}\left(\xi^2-1\right)\eta(\eta+1) \\
N_8 &= -\tfrac{1}{2}\xi(\xi-1)\left(\eta^2-1\right) \\
N_9 &= \left(\xi^2-1\right)\left(\eta^2-1\right)
\end{aligned}
$$

Figure 2.2: The 9-node bi-quadratic quadrilateral element and its shape functions

using the chain rule of partial differentiation. This results in:

$$\begin{bmatrix} \frac{\partial}{\partial x} \\ \frac{\partial}{\partial y} \end{bmatrix} = \boldsymbol{J}^{-1} \begin{bmatrix} \frac{\partial}{\partial \xi} \\ \frac{\partial}{\partial \eta} \end{bmatrix} \tag{2.36}$$

where $\boldsymbol{J}$ is the Jacobian matrix. The integrals inside the discrete weak form also need to be transformed. This can be done with the determinant of $\boldsymbol{J}$, $\det \boldsymbol{J}$ as follows:

$$\int\int dxdy = \int\int |\det \boldsymbol{J}|\, d\xi d\eta \tag{2.37}$$

What remains is to evaluate the integrals, which could be done analytically. However, for general purposes they are evaluated numerically using Gaussian quadratures rules which states that

$$\int\int f(\xi,\eta)\, d\xi d\eta = \sum_{j=1}^{n} w_j f(\xi_j,\eta_j) \tag{2.38}$$

which will give an exact result for polynomials of degree $2n-1$ or less by a suitable choice of the points (ξ_j, η_j) within the element and weights w_j. For $n=2$, these are $w_i = [1,1]$ for points $\left[-\sqrt{1/3}, \sqrt{1/3}\right]$ and for $n=3$, these are $w_i = [5/9, 8/9, 5/9]$ for points $\left[-\sqrt{3/5}, 0, \sqrt{3/5}\right]$. With the techniques explained above, the contribution of the integrals in equation 2.29 can be numerically solved for every individual element. After this is done the contribution of every element is assembled (stored in the right spot) into one big matrix. Please note that these matrices are very sparse in most cases.

The (bilinear) quadrilateral element discussed above is in most cases accurately enough, however if better solutions are needed the type of element can simply be changed into one that is more accurate while the underlying theory (and computer programme) remains the same. The most obvious way of improving the accuracy is using quadratic shape functions instead of linear ones. This means that on each side of the element one node is added. However, for reasons explained later when the finite element method is applied for Stokes flow, one extra node will be added in the centre of the element. Figure 2.2 demonstrates this 9-node bi-quadratic quadrilateral element and its shape functions.

This concludes the basic theory of the finite element method. In the next subsections it will be shown how the method can be used to compute electric fields, temperature, and Stokes flow.

2.2.2 Applied to Electrokinetic and Thermophoretic Flows

In the previous section the basic steps to solve a boundary value problem with the finite element method were explained. In this section these techniques are applied to the continuum equation for electrokinetic flows which were derived in section 2.1.2 and the thermophoretic flows derived in section 2.1.3. It was shown that in the case of microfluidics, the electric double layer generally is much smaller than any typical dimension of the microfluidic device. The consequence of this is that the effects of the electric double layer merely act as a boundary condition for the bulk flow and the electric field that needs to be simulated is only

the externally applied electric field $\vec{E}$, which is given by $\vec{E} = -\nabla\phi$, where the potential, ϕ, is the applied voltage on the electrodes and can be computed with the Poisson equation:

$$\nabla^2\phi = 0 \tag{2.39}$$

It was also shown that in the opposite case where the electric double layers do overlap considerably (e.g. in small nanofluidic devices), the flow is still directly related to the externally applied electricity field. In the case of electrophoresis the velocity of the particles or bead are also directly related to the electric field. Therefore in order to compute the flow, equation 2.39 needs to be solved. This Poisson equation was already rewritten for the finite element method in the previous section, the only difference is that the source term here disappeared. So the discrete weak form becomes:

$$\sum_{A\epsilon\lambda\setminus\lambda_D} \phi_A \int_\Omega \nabla N_B \cdot \nabla N_A d\Omega = \int_{\Gamma_N} N_B \left(\vec{n}\cdot\nabla\phi\right) d\Gamma - \sum_{A\epsilon\lambda_D} \phi_D\left(\vec{x}_A\right) \int_\Omega \nabla N_B \cdot \nabla N_A d\Omega \tag{2.40}$$

The Neumann boundary condition contains the term $(\vec{n}\cdot\nabla\phi)$, which can be considered as the amount of electric current leakage. In most microfluidic devices the walls can be assumed to be good insulators, which means that the normal component of the electric field (and the complete Neumann boundary condition) is set to zero. When equation 2.40 is rewritten to a matrix form, the result is:

$$\boldsymbol{K\phi} = \boldsymbol{f} \tag{2.41}$$

where $\boldsymbol{\phi}$ is the vector with only the unknown (internal) values of the potential and $\boldsymbol{f}$ is computed using the $\boldsymbol{K}$-matrix and the Dirichlet boundary conditions. The value of the Dirichlet boundary condition is the applied voltage to the electrode. Once the matrix is solved, the value of the potential is known in the whole device. However, in order to compute the electroosmotic flow and electrophoretic velocity of a particle or bead, the value of the electric field must be computed. In order to do this, the fact that $\vec{E} = -\nabla\phi$ can be used again. The discrete weak form of this equation is given by:

$$\sum_{A\epsilon\lambda} \vec{E}_A \int_\Omega N_B N_A d\Omega = \sum_{A\epsilon\lambda} \phi_A \int_\Omega N_A \nabla N_B d\Omega \tag{2.42}$$

This equation must be solved for the two dimensions of the electric field, which in matrix form for the x-component means that:

$$\boldsymbol{ME_x} = \boldsymbol{G_x\phi} \tag{2.43}$$

and similar for the y-component:

$$\boldsymbol{ME_y} = \boldsymbol{G_y\phi} \tag{2.44}$$

Here the $\boldsymbol{G}$-matrix is split into a matrix acting as the gradient in x-direction and y-direction. Once the electric field is computed the electroosmotic velocity can simply be computed by equation 2.8 or equation 2.10, while the electrophoretic velocity can be computed by equation 2.11. Di-electrophoresis can computed very similar to the way how the electric field is computed. The dielectrophoretic velocity for a spherical particle is given by equation 2.16 and is dependent on the gradient of the electric field intensity. Therefore, first the electric field intensity is computed and then the $\boldsymbol{G_x}$, $\boldsymbol{G_y}$, and $\boldsymbol{M}$ matrices can be used again to compute the gradient in x and y-direction.

In the case of thermophoretic flows, it was shown that the flow is caused by a temperature gradient. Although there are several ways to simulate temperature inside a continuum domain, the most straightforward one is the heat equation, giving the temperature distribution through time:

$$\frac{\partial T}{\partial t} = \alpha\nabla^2 T \tag{2.45}$$

where α is called the thermal diffusivity and T is the temperature as a function of time and x, y, and z. In the case of a steady state, this equation simply reduces to:

$$\nabla^2 T = 0 \tag{2.46}$$

and is solved exactly the same as the Poisson equation for the electric potential. The boundary conditions can be a certain (wall) temperature or a temperature gradient.

2.2.3 Applied to Stokes Flow

In the last section it was shown that variables like the temperature and electric field can easily be obtained using the finite element method. The same can be said about obtaining the velocity field and pressure field for Stokes flow, but there is one issue that need to be taken care of first. The flow in all cases discussed hereafter are considered to be incompressible and it is the incompressibility condition that is the problem. The continuity equation, equation 2.1, takes care of the incompressibility, but only puts a constraint on the velocity field and not the pressure, which is now a variable without any constitutive equation to describe it. On the other hand, in the Stokes equation, equation 2.6, the pressure is an added variable that adjust itself (instantaneously) in order to satisfy the incompressibility condition. Clearly this is not unique to the finite element method and whatever discretisation or numerical method is chosen, in the end the system that needs to be solved will look similar to:

$$\begin{bmatrix} \boldsymbol{K} & \boldsymbol{G} \\ \boldsymbol{G^T} & \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{p} \end{bmatrix} = \begin{bmatrix} \boldsymbol{f} \\ \boldsymbol{h} \end{bmatrix} \tag{2.47}$$

where $\boldsymbol{K}$ is a square matrix, $\boldsymbol{G}$ a rectangular matrix, and $\boldsymbol{u}$, $\boldsymbol{p}$, $\boldsymbol{f}$, and $\boldsymbol{h}$ have the corresponding dimensions. The problem of course is that this matrix system can become singular. If standard discretisation techniques are used, the computed pressure will be highly oscillatory or simply unknown. In finite difference methods this can be solved by using staggered grids, where the variables of pressure and velocity are computed on different nodes. Something similar can be used for the finite element method, where the velocity and pressure are computed using different (or mixed) elements. However, not every combination of elements for the velocity and pressure variables are suitable for a stable solution. The mixed elements that are suitable for finding a stable solution do follow the so-called Ladyzhenskaya-Babuška-Brezzi (LBB) compatibility condition [12, 28, 116]. Basically the LBB condition states that velocity and pressure spaces cannot be chosen arbitrarily because a link between them is necessary. This link ensures that the matrix system is not singular. Luckily, several combinations are available, and the one that is used here is the *Taylor-Hood element* [187]. In that case the velocity field is solved on a 9-node bi-quadratic quadrilateral element, which was already introduced above. The pressure on the other hand is solved using the bi-linear quadrilateral element, where the local node numbers 1, 2, 3, and 4 (see figure 2.2) are shared among the two element types.

Next the strong form is converted into weak form. The system of equations and boundary conditions in strong form can generally be given by:

$$\begin{aligned} -\nabla p + \mu_f \nabla^2 \vec{u} + \vec{f_b} &= 0 && \text{in } \Omega \\ \nabla \cdot \vec{u} &= 0 && \text{in } \Omega \\ \vec{u} &= \vec{u}_D && \text{on } \Gamma_D \\ -p\vec{n} + \mu_f \left(\vec{n} \cdot \nabla\right) \vec{u} &= \vec{t} && \text{on } \Gamma_N \end{aligned} \tag{2.48}$$

where the last equations are the boundary conditions on the Dirichlet and Neumann boundary. The weak form of the Stokes equation requires the introduction of weighting functions and trial solutions for the velocity field and the pressure field but also needs to take care of the incompressibility condition. Similar to what was defined before, for the velocity field the weighting functions, now noted by $\vec{w}$, vanish on the Dirichlet part of the boundary, Γ_D, while the trial solutions satisfy the boundary conditions. The weighting functions and trial solutions of the pressure field can be defined similar. However, it is possible to solve the system without specifying any pressure boundary condition, in which case the pressure is defined up to a constant. The incompressibility condition requires the definition of an extra set of weighting functions, defined as q. With the definition of these functions, the strong form of the system of equations can be converted into the weak form. The pressure gradient term and the viscous term can be rewritten using the divergence theorem, while the other term is very straightforward. If all the terms are added up and the Neumann boundary conditions are taken into account, the resulting equation in weak form becomes [52]:

$$-\int_\Omega p \left(\nabla \cdot \vec{w}\right) d\Omega + \mu_f \int_\Omega \nabla \vec{w} : \nabla \vec{u} d\Omega - \int_\Omega q \left(\nabla \cdot \vec{u}\right) d\Omega = \int_\Omega \vec{w} \vec{f_b} d\Omega + \int_{\Gamma_N} \vec{w} \vec{t} d\Gamma \tag{2.49}$$

The discretisation of this equation requires the introduction of the approximate solutions of the velocity components v_i^h and pressure p^h, as well as their associated weighting functions

w_i^h and q^h. The way to proceed is similar to what was done before, however the results are somewhat messier because of the vector components. For the details the reader is referred to [52], but in the end the matrix system is exactly the same as equation 2.47. In this system the (square) $\boldsymbol{K}$-matrix is the viscosity matrix containing the discretised viscosity term and because of the velocity components has the size of the number of element nodes used for the velocity times the number of dimensions. The $\boldsymbol{G}$-matrix, similar to what was the case before, is the discrete gradient operator and contains the pressure gradient term. This rectangular matrix has the dimension of the viscosity matrix in one dimension times the number of element nodes used for the pressure in the other dimension. The $\boldsymbol{G^T}$-matrix acts as the discrete divergence operator and contains the incompressibility condition. The vector $\boldsymbol{f}$ contains the body force term, the Neumann boundary condition term, but also the effect of the velocity, $\vec{u}_D$, on the Dirichlet boundary originating from the viscosity matrix. The vector $\boldsymbol{h}$ contains the effect of the velocity, $\vec{u}_D$ originating from the compressibility condition.

The way how the matrix system is solved needs some more attention. The system of equations is (very) sparse and therefore any solver for sparse systems should find the solution. However, for efficient solving of the system and memory requirements of the solver, it is also important that all the values are near the diagonal of the matrix. This can be accomplished by reordering the nodes where the global node numbers are changed in such a way that in the end, when the final matrix is assembled from the separate element, all values are as near to the diagonal of the matrix as possible. The reordering method used here is the Reverse Cuthill-McKee algorithm [44]. If this method is used, even systems with up to $200,000$ nodes can still be solved easily with any direct (non-iterative) method on a modern desktop computer.

This concludes this section where the basic theory of the finite element method was explained and where it was shown how the method can be used to discretise the continuum equations to compute the electric field, temperature, velocity field, and the pressure inside any micro-sized or nano-sized device. In the next section a different kind of discretisation is explained. Here, instead of continuum equations being discretised, the system itself is discrete in nature.

2.3 Modelling Molecules: Molecular Dynamics (MD)

In the last section the concept of a continuum was used to derive (continuum) differential equations that can describe continuum (e.g. pressure driven) flows, electrokinetic flows, and thermophoretic flows very well. This is despite that all these flows originate from molecules or atoms interacting with each other, for example, electrokinetic flows where the charge of individual atoms, the charged wall and the Brownian motion result in the electric double layer. However, there may be cases where the continuum hypothesis fails, for example when the size of the device of interest is very small or when the behaviour of one or several molecules make a huge difference, e.g. when studying the effects near (molecular) walls. In that case one would like to model the atoms or molecules themselves rather than the continuum behaviour of many of them. On the other hand, with a modelling technique that can do this, it is also interesting to study how many molecules can make a continuum or how continuum parameters like viscosity can be determined from them. This section will introduce one of such modelling techniques called *Molecular Dynamics* (MD) where the motion of individual atoms and molecules is simulated. A more detailed introduction can be found in for example [8, 63, 156].

First the basic framework is explained in section 2.3.1, followed by section 2.3.2 where some standard techniques which are prerequisites for any simulation are explained. However, for more practical simulations more advanced techniques are needed. These are given in section 2.3.3. This section ends with two subsections where two commonly used atomic systems are explained. In 2.3.5 the parameters for the simulation of inert materials like argon are given, while in section 2.3.6 it is explained how a polar liquid like water can be simulated.

2.3.1 The Essentials of Molecular Dynamics

Molecular dynamics is a numerical technique where the motion of individual atoms and molecules is simulated. However, a system of interacting atoms actually consist of nuclei and electrons that interact with each other, or more in general, quantum physics should be employed in the form of the Schrödinger equation that describes how the quantum state (or wave function) of the system changes in time. This is definitely possible for simple systems, but for more realistic problems approximation schemes have to be employed. Molecular dynamics basically employs two very reasonable approximations in order to be of practical use. The first approximation is called the *Born-Oppenheimer approximation* [25], which takes into account that the nuclei are much heavier than the electrons and that the electrons move a lot faster than the nuclei. This fact can be used to rewrite the Schrödinger equation into two separate equations. The first equation takes into account that from the electron's perspective the nuclei can be regarded as fixed, and because of that a term enters the equation that is only dependent on the coordinates of the nuclei. This term is called the *interatomic potential*. In the second equation this term is also present, and the rest of the equation does not contain any dependency on the coordinates of the electrons. All effects due to electron movement are incorporated into the interatomic potential. Hence, this implies that once the interatomic potential is known, the motion of the nuclei can be computed. The second approximation that is employed in MD is the fact that under the right circumstances (i.e.. no quantum effects), the Schrödinger equation can effectively be replaced with Newton's laws of motion. This means that in MD the motion of the nuclei (or atoms) are described by classical mechanics.

The first step in order to do a molecular dynamics simulation is to specify the interatomic or interparticle potential, i.e. a certain function that depends on the atom coordinates. Basically there are two ways of doing this. The first is to use empirical potentials, while the other is to solve the electronic structure problem, defined by the two equations mentioned in the introduction. The latter technique is called *ab initio* molecular dynamics, or the Car-Parrinello method named after Car and Parrinello [33] who developed the method. The strength of this method is that the potential is computed rather than empirically obtained. However, the disadvantage of this method is that it requires much larger computational effort. In order to do computations of reasonably sized MD systems (> 10000 atoms) there is the need for empirical potentials. Obviously the choice of the used empirical potential depends on the material that is modeled. Arguably the most popular potential used is the Lennard-Jones 12-6 potential, given in equation 2.50, after Sir John Edward Lennard-Jones. This potential is especially useful to describe the interaction between noble gas atoms and non polar molecules. The potential is [121]:

$$U^{LJ}(r_{ij}) = 4\epsilon \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - \left(\frac{\sigma}{r_{ij}} \right)^{6} \right] \tag{2.50}$$

where ϵ defines the depth of the potential well or the interaction strength with units of energy ($J = Nm$), σ is the distance (m) at which the interparticle potential is zero, and r_{ij} is the distance between the atom i and atom j. The first term inside the potential, with power r^{-12}, defines a repulsive term similar to Pauli repulsion, which acts as electrostatic repulsion between the electron clouds of two atoms when they are close to each other. The second term, with power r^{-6}, describes a long-range attraction, due to van der Waals forces. Figure 2.3 shows the interaction energy versus the interatomic distance. Note that the potential is an approximation, especially because the repulsive term should depend exponentially on the distance. However, the exponent 12 is chosen simply because it is easy and efficient to compute it from the exponent 6 of the attractive term.

Although the interparticle potential describes how the atoms interact with each other, it is the resulting force that affects the atoms' motion. By definition, a force field can be determined by taking minus the vector gradient of the potential, $\vec{f} = -\nabla U$. The force field resulting from the potential 2.50 is:

$$\vec{f}_{ij} = \frac{48\epsilon}{r_{ij}^2} \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - \frac{1}{2} \left(\frac{\sigma}{r_{ij}} \right)^{6} \right] \vec{r}_{ij} \tag{2.51}$$

where $\vec{f}_{ij}$ is the vector defining the force on atom i, resulting from atom j, $\vec{r}_{ij}$ is the vector defining the distance of atom i and j in every separate direction (x, y, z), while r_{ij} is the

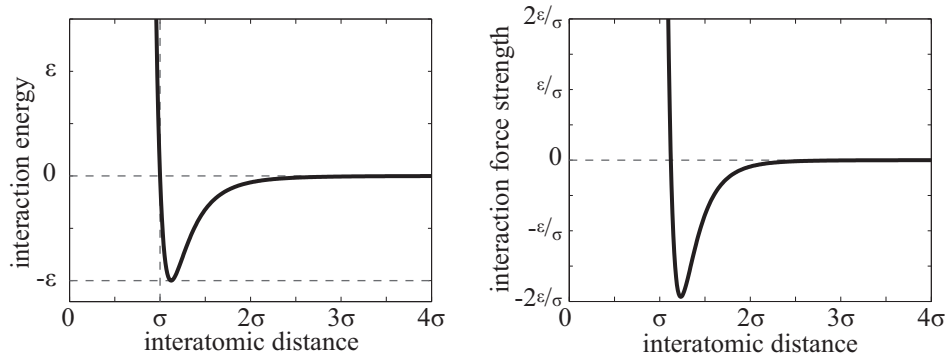

Figure 2.3: The Lennard-Jones 12-6 potential (left) and resulting force (right). Here ϵ defines the depth of the potential well, σ is the distance at which the interparticle potential is zero. At short range the potential is repulsive, while at long-range the potential is attractive.

total distance between atom i and j. Figure 2.3 shows this force as a function of interatomic distance.

Under the conditions mentioned before, the motion of the atoms can simply be computed using Newton's laws of motion. Applying Newton's second law to one atom results into:

$$m_i \frac{d^2 \vec{r}_{ij}}{dt^2} = \vec{f}_{ij} \tag{2.52}$$

where m_i is the mass of atom i. To obtain the motion of all atoms in time, equation 2.52 must be integrated with respect to time for all i.

There are several obvious factors that dictate whether a certain integration method is a good one or a bad one to integrate equation 2.52 for MD systems, but there are also some less obvious factors [63]. For example the speed of the method is not at all important, because in most MD simulations most time (>85%) is spent on the identification of interacting atoms and the computation of the interparticle forces. On the other hand it would seem important to be able to use large time steps, exactly for the same reason mentioned. However, energy conservation is very important, although two kinds of energy conservation must be distinguished; short term (<100 time steps) and long term (millions of time steps). In general, high-order algorithms with respect to the time step have very good short term energy conservation, but tend to drift in the long term. Because a MD simulation basically computes the trajectory of atoms, one could think that accurately predicting the positions of all atoms in time is the most important factor any algorithm should fulfill, however such an algorithm does not exist [63]. Luckily that is not the aim of a MD simulation; the aim is to collect meaningful statistical (average) information about the system and fortunately that is possible with very simple algorithms as long as the time integration algorithm is symplectic [168]. Then finally the algorithm should be time reversible, as the equations of motion also are.

The *Verlet algorithm* [199, 200] is a very good algorithm for integration of the equations of motion in a MD simulation and definitely the most commonly used time integration algorithm. The derivation starts with two Taylor expansions of the coordinates of an atom, around time $t + \Delta t$ and $t - \Delta t$. The summation of these two expansion is:

$$r_i(t + \Delta t) + r_i(t - \Delta t) = 2r_i(t) + \frac{\partial^2 r_i(t)}{\partial t^2} \Delta t^2 + \mathcal{O}\left(\Delta t^4\right) \tag{2.53}$$

which after rearranging, using equation 2.52 and dropping the fourth-order and higher-order terms in Δt can be rewritten to:

$$r_i(t + \Delta t) \approx 2r_i(t) - r_i(t - \Delta t) + \frac{f_{i,j}}{m_i} \Delta t^2 \tag{2.54}$$

Note that the value of the velocities of the atoms are not required to compute the new position of the atoms. However, most of the time the value of the velocity of the atoms is still required, e.g. for the computation of the kinetic energy. In this case the value of the

velocity, u, is easily obtained from the central difference,

$$u(t) = \frac{r(t+\Delta t) - r(t-\Delta t)}{2\Delta t} + \mathcal{O}\left(\Delta t^2\right) \tag{2.55}$$

The second-order accuracy is sufficient for most cases. However, if needed, more advanced integration schemes can be used [63]. Instead of the Verlet-algorithm, another frequently-used algorithm is the so-called *velocity Verlet-algorithm* [185]. The new position of the atom is then computed with:

$$r_i(t+\Delta t) \approx r_i(t) + u(t)\,\Delta t + \frac{f_{i,j}}{2m_i}\Delta t^2 \tag{2.56}$$

and the new velocity of the atom is computed with

$$u_i(t+\Delta t) \approx u_i(t) + \frac{f_{i,j}(t+\Delta t) + f_{i,j}(t)}{2m_i}\Delta t \tag{2.57}$$

In this algorithm the new velocity can only be computed once the new position of the atom is known and, resulting from that, the new force on the atom. However, it can be shown that the velocity Verlet-algorithm is exactly the same as the "normal" Verlet-algorithm. However, in this way both the position and velocity are known at the same time. Although the Verlet algorithm is not accurate for long time steps, the main reason the Verlet algorithm is frequently used, is because of its little long-term energy drift compared to other algorithms [63]. Also, the time step is severly limited anyway, because of the very high forces involved in the MD simulation when two atoms come close to eachother. Furthermore, the size of the time step will also be dependent on which atoms are simulated, and will be discussed in section 2.3.5 and 2.3.6. The simulation code used in this thesis is using the velocity Verlet-algorithm.

This concludes this section, where the bare minimum computational requirements were described for a MD simulation. Any MD simulation code first computes the interparticle forces using the current positions of all atoms and from those forces, the new positions (and velocities) are computed, and the loop continues. However, to start a MD simulation some initial conditions and boundary conditions need to be specified. Further, to simulate relevant situations, control those situations, post-process them, and compute them efficiently, a whole range of other techniques is required. The next section describes these.

2.3.2 Standard Techniques for Molecular Dynamics

All simulations need boundary conditions, with MD simulations not being any different. However, for a continuum boundary conditions are normally specified for a certain differential equation. In that case a boundary condition can be a certain velocity, pressure, temperature, or even a certain gradient of a property. Although in MD simulations Newton's laws of motion are solved, the boundary conditions must be applied to the atoms or particles. Properties like velocity or pressure only make sense when many atoms are averaged over many time steps, however the boundary conditions for the atoms need to be applied at every time step. Therefore, specifying a certain (continuum) velocity to a MD system is not straightforward and will be discussed in Chapter 3. Luckily, several boundary conditions are easy to implement and are very useful for obtaining relevant results with MD simulations.

Before the actual boundary conditions are discussed, it is first investigated to what the boundary conditions need to be applied to. Figure 2.4 shows the simulation domain, the so-called MD system box. This is a box with a certain volume that contains a certain number of atoms that all interact with each other. All atoms also have a certain velocity, and the total system has a certain amount of total energy. The MD system box used throughout this thesis is three-dimensional, but one or two-dimensional systems can be simulated with the same principles. The MD system box has in total 6 boundaries and for all these boundaries a boundary condition needs to be specified. Because MD simulations are progressing in time, an initial condition is also needed. This initial condition must be applied to both the boundaries and the interior of the MD system box. This means that on each atom an initial condition must be applied.

The simplest boundary condition to implement is a periodic boundary condition. In MD simulations a periodic boundary means that if an atom for example crosses one system

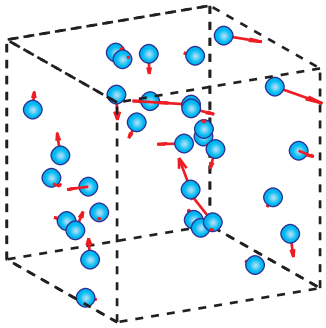

Figure 2.4: The MD system box. A three dimensional box with a certain volume and number of atoms. The system has a certain amount of energy, and it has six boundaries.

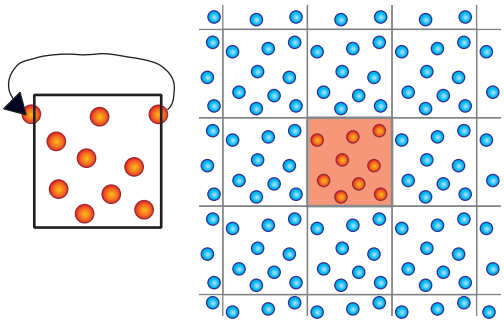

Figure 2.5: Periodic boundary conditions in MD. The central MD system now interacts with an infinite number of copies of itself.

boundary, its position is changed to that where it just passed the opposite system boundary in the box. The left diagram in figure 2.5 shows this occurrence. This concept is very simple, however one must keep in mind that when simulating periodic boundary conditions, actually an infinite number of periodic MD systems is simulated. This is illustrated in the right diagram of figure 2.5. In other words, the atoms inside our MD box interact with the imaginary atoms inside the other MD boxes, which basically are copies of themselves. Clearly a trick must be used, else the MD simulation would consume infinite time.

One particular method is frequently employed to simplify the computation of the interactions, which is the use of a cutoff radius. The idea is simple, especially when a Lennard-Jones potential is applied. Inspecting Figure 2.3 reveals that when the atoms are in close range ($r_{ij} < \sigma$) or medium range ($r_{ij} < 2\sigma$), the interaction energy is large. However, when the atoms are far away from each other ($r_{ij} > 2\sigma$), the interaction energy vanishes and rapidly becomes many times smaller than the interaction at closer range. The cutoff method makes use of this fact by only considering the interactions between atoms that are within the cutoff radius and thus significantly reduces the computational effort of MD. However, choosing the right value of the cutoff radius is still a challenge. Commonly used cutoff radii in literature are either $r_c = 2.3\sigma$ or $r_c = 2.5\sigma$, where the interaction energy has dropped to only 1.6% and 0.8% of the maximum possible attraction value, respectively. Figure 2.6 shows this cutoff radius for the Lennard-Jones potential. For every atom, only the interaction between all atoms inside the surrounding sphere of radius r_c need to be computed. In most cases, simulations using these cutoff radii give close to satisfactory results while keeping the computation time within acceptable limits. However, when longer range attractive forces are important, which is the case when simulating surface tension or electric forces, one must use larger cutoff radii and/or use additional techniques. These techniques are explained in section 2.3.4.

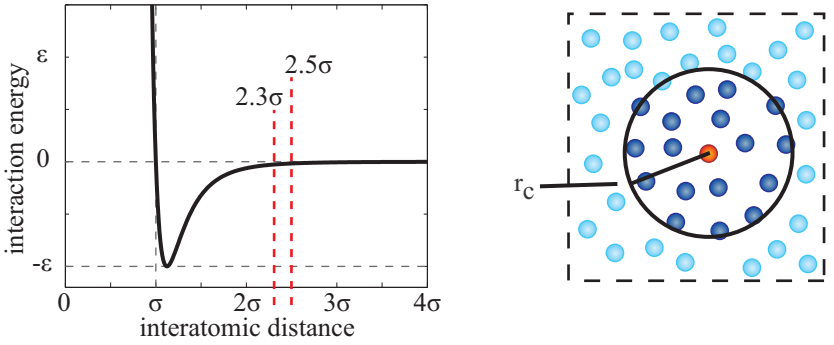

Figure 2.6: To optimise the computation a cutoff radius is used. The interparticle interaction for the central atom is only computed within this radius.

Periodic boundary conditions are ideal for simulating bulk properties of the fluid. However, another interesting problem to simulate is the behaviour of liquids near walls or objects. In continuum computations this boundary condition is specified by, for example, the no-slip boundary condition (velocity is zero at the boundary), or more general, a certain velocity at the boundary. As explained before, specifying a certain (continuum) velocity is not straightforward in MD, because this is an average property over many atoms and many time steps. Luckily the solution is very elegant; instead of specifying a certain artificial boundary condition, the molecular structure and behaviour of the wall can be modelled with MD as well. Although in theory a complete modelling of the solid wall with specialised potentials for certain metals or crystals is possible. However, the required time step to integrate the very strong intermolecular forces and the required time step to capture the vibrations of the atoms inside the wall, which are generally smaller than those required for the MD simulation of the atoms for the liquid, makes this solution computationally expensive. Luckily several simplified models can be used that still produce realistic physical behaviour, which is shown in Section 2.4. The simplest model is where all the interactions of the wall working on the liquid is incorporated into a single potential that only depends on the normal distance from the wall. This so-called "smooth wall" potential [177, 181, 193] can be obtained from the integrated interaction energy from all the atoms inside the wall. One of the assumptions of this potential is that the structure of the wall, parallel to the wall, is the same everywhere. Therefore this wall potential can not be used for simulating a wall with molecular imperfections or a curved wall. In that case it is still necessary to compute the atomic structure of the wall in a simplified manner. This method is known as an "atomistic wall" where it is necessary to compute the interaction with each of the wall atoms individually. Figure 2.7 shows the case of atoms composing a liquid in the vicinity of an imperfection in the solid (atomistic) wall. The strength of this method is that the wall atoms do not need to be identical. In other words, they can for example represent atoms of different types of metals. This can be controlled by specifying the interaction energy of each atom and its interaction distance with the atoms in the fluid, which are ϵ_{wf} and σ_{wf}, respectively. However, as explained before, to simulate actual solid atoms, the time step needs to be made smaller than required for the fluid atoms and possibly many more atoms need to be modeled. Therefore to allow efficient computations, some simplifications have to be made. The first of such simplifications is to constrain the atoms to their lattice site with springs [176, 189, 194], thus simulating thermal motion only. It is also possible to actually fix the atoms to their lattice site [85, 216]. In this case a direct thermal motion exchange is not simulated, but other physical properties of solid-fluid interaction, like wetting, can still be modelled. Finally, another wall boundary condition can be applied to a MD system boundary, i.e. the specular wall. A specular wall changes the sign of the velocity component of the atom perpendicular to this wall. This means that no interparticle interactions are modelled with the wall. This boundary condition is mostly suitable for gas atoms, because in that case the interparticle interactions are less important. For atoms representing a liquid the applicability is therefore limited, but as long as the boundary is far away from the simulation area of interest, it is an efficient way of keeping all atoms inside a MD system box.

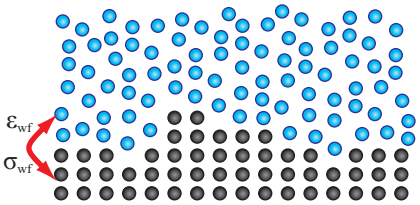

Figure 2.7: An atomistic wall and the interaction between the wall and atoms in a liquid which can be controlled by changing ϵ_{wf} and σ_{wf} to simulate different material and surface properties.

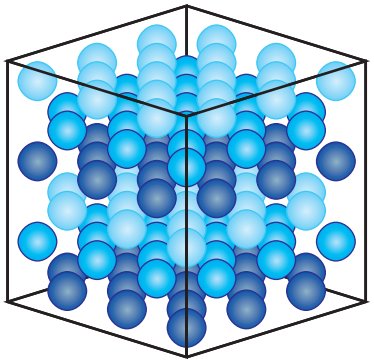

Figure 2.8: A typical initial solution for an MD system. All atoms are ordered in an FCC lattice and are given initial velocities that obey Maxwell's distribution at a given temperature. The simulation starts by "melting" the system. The different colours of the atoms indicate the different layers of the lattice.

As mentioned before, besides the boundary conditions, any simulation progressing in time also needs an initial condition. Similar to the case of a steady state continuum, to improve the speed of convergence, one would like an initial condition that is in a state close to the finally desired solution. However, unlike in continuum where only several variables need to be determined, in MD many different positions and velocities of the atoms correspond to basically the same state of the system (total energy, temperature, etc.). Furthermore, the equilibrium properties of the system should not depend on the initial conditions, therefore the initial condition can be determined by convenience. If the solid state of a certain molecular material is simulated, it makes sense to initialise the system in its known crystal structure. For example, for metals one should choose either the bcc (body-centred cubic), the fcc (face-centred cubic), or the hcp (hexagonal close packed) crystal structure. The distance between the lattices is determined by the desired density of the substance. Similarly, to model the liquid state of a substance, one also initialises the system in one of the crystal structures, and subsequently the system will "melt" once the simulation is started [63]. This is because, at the temperature and density of the liquid state, the solid state is not thermodynamically stable. Figure 2.8 shows the initial condition of a Lennard-Jones fluid at density ρ and the resulting lattice distance. The initial condition is not complete yet with only the initial position of the atoms specified. To complete the initial condition, the velocity of every atom needs to be specified. The most convenient way to do this, is to assign to the atoms the velocities according to a Maxwellian distribution at a given temperature. However, as the system is melting, the actual temperature of the system will change, so that property is not critical.

Now the initial conditions are defined, the boundary conditions have been specified, and the interparticle potential is chosen, the basic MD simulation can commence. First the interparticle forces are computed from the current positions of the atoms. These forces are used to advance the positions in time and if needed, the (instantaneous) velocities of the atoms are computed. Then the time iteration step is completed and the simulation

is continued to the next iteration. However, the purpose of a MD simulation is to obtain macroscopic properties of a certain atomic substance. In other words, the average property of many atoms over a long enough time is wanted. The collection of meaningful macroscopic properties such as pressure and temperature from MD simulations is commonly known as "binning". Figure 2.9 demonstrates this. At certain times during the simulation, which could be every n-th time step, all known (instantaneous) simulation parameters of all atoms whose position is inside the bin are collected. For example, to obtain the (local) density of the fluid inside the bin, one simply adds up the masses of the atoms inside the bin, and divides that number by the volume of the bin. This step yields the average wanted property in space. By taking the average of this property for the same bin through time, the property is both averaged in space and time, and results in the macroscopic density. Other commonly binned properties are (macroscopic) velocity, (macroscopic) temperature, and (macroscopic) pressure.

For the (macroscopic) velocity, the bin simply collects the average instantaneous velocity or momenta of all the atoms inside the bin through time. Similarly, the (macroscopic) temperature is obtained by averaging the so-called instantaneous temperature inside the bin through time. The instantaneous temperature however is not a simulation variable that is readily available, like position or momentum, and therefore must be determined from the simulation variables somehow. This is the field of classical statistical mechanics, and the theory that translates the simulation variables into a temperature is called the equipartition theorem [204]. This theorem states that in thermal equilibrium each microscopic degree of freedom contains an amount $\frac{1}{2}k_bT$ of thermal energy associated with it. In case of N atoms with s degrees of freedom, this means that the total internal energy is given by:

$$U = N\frac{s}{2}k_BT \tag{2.58}$$

In case the degrees of freedom are only the three translational degrees of freedom, this energy corresponds to the three components of velocity, i.e. the total kinetic energy. Rewriting this for the temperature results into

$$T_{inst} = \frac{1}{3k_BN}\sum_{i=1}^{N} m_i\left(v_{i,x}^2 + v_{i,y}^2 + v_{i,z}^2\right) \tag{2.59}$$

When more complex molecules are modelled, the number of degrees of freedom changes. In that case there are also rotational degrees of freedom and vibrational degrees of freedom, which for example refer to the bonds between individual atoms. This change in the degrees of freedom will be further discussed in Section 2.3.6, where it will be explained how a more complex molecule, like water, can be represented in a MD simulation.

Likewise, the (macroscopic) pressure must also be determined from the simulation variables. There are several ways to do this [184], but the most common one is based on the virial equation for the pressure. The virial equation for the pressure describes the stresses in terms of the positions, velocities, and resulting interparticle forces. The components (α and β) of the stress tensor are given by [96]:

$$\Sigma_{\alpha\beta} = \frac{1}{V}\left[\sum_{j}^{N} m_j v_{j,\alpha} v_{j,\beta} + \frac{1}{2}\sum_{i,i\neq j}^{N}\sum_{j}^{N} r_{ij,\alpha} f_{ij,\beta}\right] \tag{2.60}$$

where $v_{j,\alpha}$ represents the α component of the velocity of atom j. From a practical point of view, because the pressure is obtained using the interparticle forces, the value of the product in the second sum in 2.60 is determined at the same time when the interparticle force is determined. In other words, this value is determined inside the "force loop", and at the end of this loop every atom has its contribution to the total stress associated to it. The final stress is obtained by adding these contributions.

This concludes the section on the basic techniques required for a molecular dynamics simulation. However, for more realistic simulation several more advanced techniques are required. For example, a simulation where the temperature needs to be controlled, or where long-range force are important. Also, there are several techniques available that speed up the force computation, which is the most time-consuming part of any MD code. These subjects will be discussed in the next section.

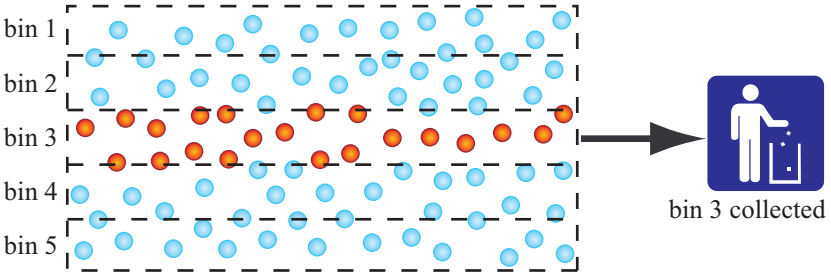

Figure 2.9: In order to obtain meaningful macroscopic or averaged properties from a MD system, like density, the instant values of the simulation are collected or "binned" every time step.

2.3.3 Advanced Techniques for Molecular Dynamics

Although the techniques and conditions explained in the previous section are sufficient to perform an MD simulation, one will find out that the number of atoms that can be simulated is severely limited by the available computational resources (i.e. memory and CPU speed). Furthermore, it is sometimes necessary to do a simulation at a certain temperature or a simulation in presence of temperature gradients. This section discusses several techniques that are required to enable such simulations.

2.3.3.1 Optimising the Force Loop

As mentioned in the previous section, in a MD simulation the interactions between all atoms is computed. So when the MD domain contains N atoms, in total $N(N-1)$ interactions must be computed. However, as long as the interparticle potential allows it, only short-range interaction can be considered to render a more efficient computation. This is accomplished by using a cutoff radius r_c as explained in Section 2.3.2. In the three-dimensional case this means that only the interactions between the atoms within a sphere of radius r_c need to be considered. This means that for a single atom, on average, the interaction between $n_{rc} = \frac{4}{3}\pi\rho_n r_c^3$ atoms need to be determined, where ρ_n is the number density. So, in total there are now only Nn_{rc} interactions to be determined. This is a substantial reduction in computational effort in comparison to the full evaluation when the number of atoms is very large. Figure 2.10 shows the "force loop" in pseudo-code. The shaded area in the code is only enabled when a cutoff radius is used. So, by using a cutoff radius, the actual value of the force on atom i between the atoms does not have to be determined for every atom j. This obviously saves some computation time, however the distance between atom i and all other atoms in the MD system still needs to be determined before it can be concluded whether it is within the cutoff radius, i.e. it still requires $N(N-1)$ operations. Luckily, for this problem there are some very simple solutions.

The goal is to minimise the (needless) amount of times the distance between two atoms is checked that are far away from each other. The first very simple approach is to divide the MD system into cells of a given size. This organisation of every atom in a certain cell takes only N operations. After this step is completed, every atom knows to which cell it belongs and every cell knows how many and which atoms are within that cell. The strength of this method, which is called the *cell list method* or *linked cell method* [8], is that now only the distance of an atom in one cell and between all atoms in the surrounding cells needs to be considered. This then strongly reduces the amount of distance checks. The cell size has not yet been specified. The most commonly used value for the cell size is actually the cutoff radius or a slightly larger value. Why this value is chosen, is explained next. Figure 2.11 shows an atom that is belonging to a certain cell (the central shaded area) which size is equal to the cutoff radius. The thick black circle shows the spherical radius of the cutoff radius in which all the interactions need to be determined. It can be clearly seen that in this case only the adjacent cells need to be checked for the interatomic distances. In this case, in three dimensions, in total 27 cells need to be considered in order to obtain all

```
for each atom i (1 to N)
{
   for each atom j except i
   {
      rij = distance between atom i and j
      if (rij < cuttoff radius)
      {
         compute interparticle force
         assign force from atom j to atom i
      }
   }
}
```

Figure 2.10: The most naive version of a MD "force loop" in pseudo code. The shaded area is the extra part of code when performing an MD simulation with a cutoff radius. Although this shows how easy the principle of a MD simulation is, in this form it is not very efficient.

interactions for a single atom. This means that for one atom, on average, the interaction between $n_{cl} = 27\rho_n r_c^3$ atoms need to be determined. Note that if the third law of Newton can be used, i.e. $f_{ij} = f_{ji}$, then the number of interactions can be reduced to $13 + 1 = 14$ cells. Nevertheless, in total there are now Nn_{cl} distance checks needed instead of $N(N-1)$, which means for reasonable $N \sim \mathcal{O}(10^3)$ a huge reduction in computational effort. The cell list method eliminates many needless computations of the distance between atoms, however the number of checks, n_{cl}, is still about 6 times higher than the maximum needed checks, n_{rc}. This can be optimised by choosing the proper size of the cell, which will be discussed later. However, first another simple but effective method to speed-up the interparticle force computation is discussed.

As mentioned before, typical time steps in MD simulations are very small because it is necessary to integrate the very steep interparticle force function. This means that the distance the atoms travel in one time step does not change considerably, and as a result the interaction between basically almost the same set of atoms need to be determined over several time steps. This can be utilised by making a list of all the atoms that are within the cutoff sphere with a radius r_c plus all atoms with a position inside a slightly larger sphere with a radius $r_c + r_{sh}$. Here, r_{sh} is the radius of the shell surrounding the cutoff sphere. The list containing these atoms is called the *neighbourhood list* [199, 200]. Figure

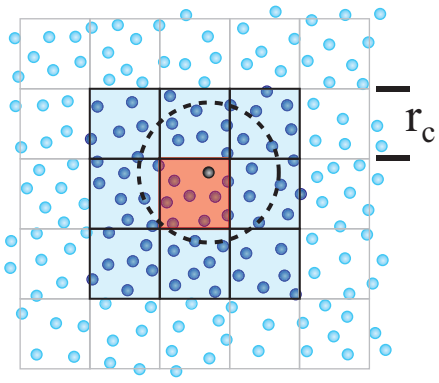

Figure 2.11: The cell list method with cell size r_c can be used to optimise the compuation of the interparticle interactions by onlt considering the central cell and only its directly surrounding cells.

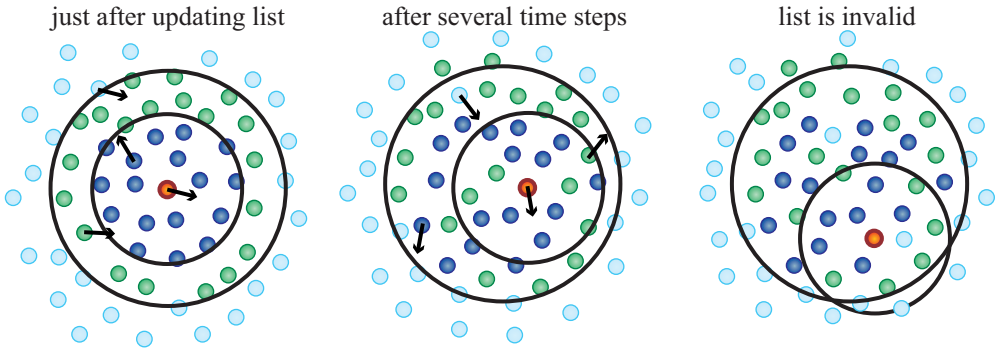

Figure 2.12: The neighbourhood list method can be used to prevent the computational effort of creating a list of all possible interactions with one atom every time step. Now the (cell)list only needs to be updated once one of the interacting atoms leaves the outer shell.

2.12 illustrates this method. The strength of this method is that once this list is compiled, no new list is needed for every time step, but only when new atoms enter the inner cutoff sphere. Further, only $n_{nl} = \frac{4}{3}\pi\rho\left(r_c + r_{sh}\right)^3$ distance checks need to be computed, which under most circumstances results in 5 times fewer checks in comparison with a cell list method. However, the big disadvantage of this method is the actual compilation of this list. Finding all atoms within the neighbourhood cutoff radius $(r_c + r_{sh})$ still requires $N\left(N-1\right)$ operations, although it is not needed for every time step.

The cell list method has the advantage that it is independent of the size of the MD domain and that the number of distance checks scales linearly with the number of atoms, $\mathcal{O}\left(N\right)$. The neighbourhood list method has the advantage that even fewer distance checks need to be performed, but has the big disadvantage that making the list requires $\mathcal{O}\left(N^2\right)$ operations. Therefore, a logical outcome is to combine these methods. This means that the cell list method is used to make the neighbourhood list whenever needed, which requires $n_{clnl} = 27N\rho_n\left(r_c + r_{sh}\right)^3$ distance checks. Once the neighbourhood list is created, as long as the list is valid, the force loop only needs to consider $Nn_{nl} = {}^4\!/_3\pi N\rho_n\left(r_c + r_{sh}\right)^3$ distance checks. Figure 2.13 illustrates the combination of the two methods. However, which of the three method to choose, depends on several system specific conditions, especially system size. A comparison of the three methods is for example given by Muth et al. [134].

However, As was briefly mentioned before, the cell list method can be optimised further by choosing the proper cell size. This is elaborated in Figure 2.14. The left diagram shows the situation where the cell size is equal to the cutoff radius (this is the standard way of

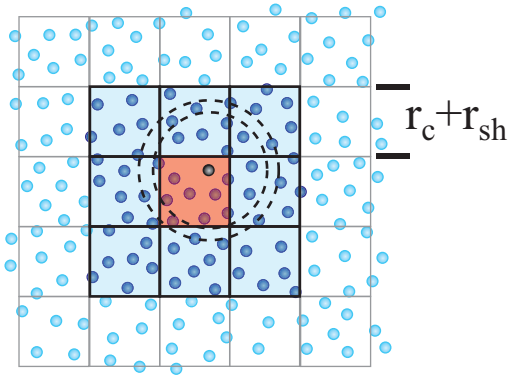

Figure 2.13: Combination of the neighbourhood list method with the cell list method.

doing this). As is evident from the diagram, with the current position of the interaction atom, it is not necessary to compute the distance with the atoms on the far left and far bottom. In other words, most atoms in the adjacent cells do not need to be considered for the interaction. The solution to this is to subdivide the MD domain into cells with a size that is half the cutoff radius. This is shown in the right diagram of Figure 2.14. Although in this case 125 surrounding cells have to be considered, or $64 + 1 = 65$ when the third law of Newton is applied (so more than the 27 or 14 in the previous case), the size of the cells are smaller. Therefore, in total only $n_{cl2} = 125 N \rho_n \left(\frac{1}{2} r_c\right)^3$ distance checks need to be performed, which is more than 1.7 times less than using the standard method. The downside of this sub-division is that more cells need to be stored and checked in the memory of the computer, which makes the duration of the loop longer. In test cases (not reported here) it was confirmed that for a small total number of atoms, $N < 2048$, the speed-up was $20 - 40\%$ compared to the standard method. For a reasonable total number of atoms, $N \approx 10^4$, the time saved was only 10% and the code even becomes slightly slower than the standard method, when N is increased beyond that ($N > 10^4$). However, for reasons explained in Chapter 3, this cell list method with cell size equal to half the cutoff radius, is used for all results obtained in this thesis.

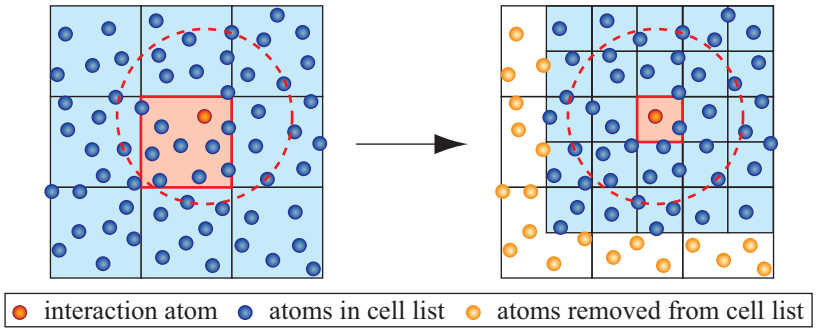

Figure 2.14: A more efficient celllist by making a celllist with $1/2 r_c$ resulting in less distance checks inside the code.

2.3.3.2 Thermostats: Controlling the Temperature

In certain cases, the purpose of the MD simulation would be to obtain the behaviour of a liquid at a given target temperature, for example, obtaining a system property like viscosity at a given temperature. However, without exactly specifying it, all simulation techniques discussed so far apply to the microcanonical ensemble or *NVE ensemble*. In other words, in every simulation the number of atoms, N, the volume of the MD domain, V, and the total energy (potential plus kinetic energy) remain constant during the simulation. For a MD simulation at constant temperature, the simulation is performed in a canonical ensemble or *NVT ensemble*. In this case heat can flow through the MD system boundaries so energy can vary, but the temperature of the MD system matches that of a (virtual) surrounding thermal bath [63]. As was discussed in Section 2.3.2, the instantaneous temperature in a MD simulation is related to the total kinetic energy. However, referring to a constant temperature simulation is not the same as saying that the instantaneous temperature or kinetic energy is constant. This is connected to the fact that the instantaneous temperature demonstrates (natural) fluctuations and by keeping the kinetic energy (perfectly) constant, these would disappear.

The simplest and most naive way of simulating a constant temperature is to actually ignore the fluctuations in the instantaneous temperature. This is opposing the facts discussed in the previous paragraph. However, it still can be used to obtain relevant results, but it will of course fail when the temperature fluctuations are important. This method is called the *velocity-rescale method*, and, as the name implies, rescales the velocities of the atoms in such a way that the total kinetic energy corresponds to the desired temperature. The kinetic energy of the MD system is changed by rescaling the velocities of all atoms by a single

scaling factor, which is equal to $\sqrt{T_{wanted}/T_{inst}}$. This scaling is applied at every time step, which consequently holds the kinetic energy of the MD system at exactly the same value. There are however weaker formulations of this approach, and one of them is the *Berendsen thermostat* [18] where the temperature is maintained by coupling the system to an external heat bath, which has a fixed (target) temperature. The velocities are still scaled at each time step, but in such a way that the rate of change of temperature is made proportional to the difference in temperature: $dT/dt = 1/\tau\,(T_{wanted} - T_{inst})$, where τ is a time scale that controls the strength of the coupling. In this case the scaling of the velocity is equal to $\sqrt{1 + \Delta t/\tau\,(T_{wanted}/T_{inst} - 1)}$, and typical values of τ in literature are taken as $50\Delta t$. Please note that when $\tau = \Delta t$, the method reduces to the velocity rescale method.

The velocity-rescale method and Berendsen thermostat are very easy to implement and generally give satisfactory results. The methods however do not exactly solve the true NVT ensemble, but fortunately several methods are available that can achieve this. The most commonly used one is called the *Nosé-Hoover thermostat* [91, 139], after Nosé who introduced the method and by Hoover who reformulated the method. The general idea is to consider the heat bath as an integral part of the system. This is done by extending the real system by the addition of an artificial variable $\tilde{s}$ associated with a mass Q and a velocity $\dot{\tilde{s}}$. Here the magnitude of Q determines the coupling between the reservoir and the real system, and consequently influences the temperature fluctuations, while $\tilde{s}$ plays the role of a time-scaling parameter. The time-scaling in the extended system ensures that although the positions of the atoms in both systems are identical, the velocities in the extended system are amplified by a factor $\tilde{s}^{-1}$ compared to the velocities in the real system. By a suitable choice of the Lagrangian for the extended system, it can be shown that although the equations of motion sample a microcanonical (NVE) ensemble in the extended system, they do sample a canonical (NVT) ensemble in the real system. However, because of the involved time-scaling, the sampling is now done at uneven time steps which can become quite complicated. Luckily, there is a way to reformulate the equations in terms of real system variables by a clever transformation [91]. The result is an equation of motion for the real system, which incorporates the coupling between the heat bath and the real system, instead of a separate rescaling.

So far the thermostats were applied to the whole MD system. In other words, the whole MD system is kept at the same temperature. However, there are cases, especially in the case of thermophoresis, were (local) temperature gradients need to be modelled. In the case of the modelling of a nano-sized channel with a temperature gradient, one could only thermostat the solid atoms in the atomistic wall. The atoms that compose the liquid near each wall will be thermostatted to the temperature of the wall by means of the solid-fluid interaction. On the other hand, it might be interesting to simulate the behaviour of a liquid in a certain temperature field. In that case the wanted temperature for the thermostat becomes a function of space. In practice this implies that the MD domain is divided in small sections where each section is thermostatted to the appropriate temperature. For convenience, and reasons explained in Chapter 3, where the coupling between MD and the continuum is explained, the subdivision is chosen to be equal to the cells used in the cell list method discussed previously. Please note that the thermostat can also be used to merely equilibrate the MD system to a certain temperature. If there are no external forces acting on the MD system, the thermostat can be switched off once the wanted temperature has been reached. From then on, the MD simulation continues to simulate the NVE-ensemble again.

2.3.4 Long-Range Forces

In this section a short overview is given of what is involved in computing long-range forces, such as electrical effects or surface tension effects, in molecular dynamics. The problem with such long-range interactions, especially Coulombic interactions that decays as ($\sim 1/r$), is that the interaction potential cannot be truncated without loosing crucial accuracy, since it is slowly decaying, while at short-range the interactions change very rapidly. Therefore, more advanced techniques need to be employed that accurately compute the long-range forces, albeit the computational effort increases considerably. Recently these long-range methods were also used for the van der Waals attraction term, i.e. the $1/r^6$ term in the Lennard-Jones potential 2.50, to more accurately describe it. Shinha et al. [175] and Bo Shi et

al. [172] showed that this improved simulation results where the long-range attraction is very important, in this case surface tension, while keeping the cutoff radius to a small value whereas normally one had to increase the cutoff value considerably to obtain similar results.

2.3.4.1 Ewald Summation

The most widely used method to compute long-range forces is *Ewald-summation* [59], where the slowly converging interacting potential is split into a short-range part and a long-range part. The key to the success of this method is that the short-range part sums quickly in real space, while the long-range part sums quickly in Fourier space. Of course, because Fourier series are used, the method only works really well for MD domains that are infinitely periodic in all dimensions. However, there are alterations to the method to work with a domain that is periodic in only 1 or 2 dimensions are periodic, which will be shown in Section 2.3.4.3. The main drawback of the method is that the method scales like $\mathcal{O}\left(N^2\right)$, where N is the number of charged atoms inside the system, or at best scales like $\mathcal{O}\left(N^{3/2}\right)$, if optimised cutoff values are used [50]. There are several ways to derive the Ewald summation equations, where the traditional method is derived using Poisson's equation for the electrostatic potential [63]. However, there are more general methods available that can also be used to get solutions even for long-range interactions other than Coulombic interactions, which will be discussed later. To illustrate the Ewald summation for Coulombic interactions, consider a cubic MD system box with length L, containing N atoms with charges q_i and positions $\vec{r}_i$. The system box has periodic boundary conditions everywhere, and the net charge of the system is zero, i.e. neutral. This means that the total electrostatic energy of the system is given by [63]:

$$E = \frac{1}{2} \sum_{\vec{n}\epsilon\mathbb{Z}^3} \sum_{i,j=1}^{N} \frac{q_i q_j}{|\vec{r}_{ij} + \vec{n}L|} \tag{2.61}$$

where $\vec{n}$ is a vector in integer space representing the discrete periodic copies of the MD domain in all three directions. The problem of this equation is that in this form it is only conditionally convergent, and therefore it is not possible to solve this equation efficiently. However, by means of Ewald summation [59], the equation can be rewritten, i.e. the original equation can be split into two equations. The first step to split the equation is achieved by assuming that each charged atom is surrounded by a diffuse charge distribution of opposite sign, which neutralises the charge of the atom, i.e. effectively screening the atom. This distribution is local and can be computed by direct computation over all atoms. The second step is to restore the original system, which is done by considering the Coulombic interaction of similarly charged distributions centred at the charged atoms, but now with the same sign as the original ions. This term exactly cancels the charge distribution before, so that the total interaction only comes from the point charges. This last term is a slowly varying function, which is best solved in reciprocal space. For the Ewald summation for Coulombic interaction, the required Fourier transforms are particularly simple when the shape of the charge distributions is chosen to be Gaussian [63]. There are however two other terms to be considered. The first one is the self interaction correction term, taking into account that the summation was over all atoms in Fourier space including the atom itself. The second one is the so-called bipolar correction, which takes into account that the modified system is surrounded by a medium with a certain effective dielectric constant ϵ_∞. In conclusion, the Ewald summation makes it possible to write the Coulombic potential into [50]:

$$U_{coul} = U^{(r)} + U^{(k)} + U^{(s)} + U^{(d)} \tag{2.62}$$

where:

$$U^{(r)} = \frac{1}{2} \sum_{\vec{m}\epsilon\mathbb{Z}^3} \sum_{i,j=1}^{N} q_i q_j \frac{erfc\left(\alpha \left|\vec{r}_{ij} + \vec{m}L\right|\right)}{\left|\vec{r}_{ij} + \vec{m}L\right|} \tag{2.63}$$

$$U^{(k)} = \frac{1}{2}\frac{1}{V} \sum_{\vec{k}\neq 0} \frac{4\pi}{\vec{k}^2} \exp\left(-\vec{k}^2/4\alpha^2\right) \left|\tilde{\rho}\left(\vec{k}\right)\right|^2 \tag{2.64}$$

are the real, fast decaying part and the slowly varying part, while the self interaction correction and bipolar correction terms are given by:

$$U^{(s)} = -\frac{\alpha}{\sqrt{\pi}} \sum_{i}^{N} q_i^2 \text{ and } U^{(d)} = \frac{2\pi}{(1+2\epsilon_\infty)\,V} \left(\sum_{i}^{N} q_i \vec{r}_i \right)^2 \tag{2.65}$$

In these equations, $\vec{k} = {}^{2\pi}/_{L}\vec{l}$, where $\vec{l} = (l_x, l_y, l_z)$, is the lattice vector in Fourier space, V, the volume of the original system box, and the Fourier transformed charge density $\tilde{\rho}\left(\vec{k}\right)$ is defined as:

$$\tilde{\rho}\left(\vec{k}\right) = \int_V \rho\left(\vec{r}\right) \exp\left(-\imath\vec{k}\cdot\vec{r}\right) d\vec{r} = \sum_{j=1}^{N} q_j \exp\left(-\imath\vec{k}\cdot\vec{r}_j\right) \tag{2.66}$$

The parameter α in these equations, which will be referred as the Ewald or splitting parameter, tunes the relative weights of the real space and the reciprocal space contributions. For computational efficiency the Fourier sum must be terminated somewhere at $\vec{k}_{max}$, and the real part uses a cutoff radius of r_{max}. The magnitude of the cutoffs can be optimised for a certain value of α, such that the accuracy of the approximated Ewald sum is the highest possible [50]. The Coulombic force field can simply be determined by $\vec{f} = -\nabla U_{coul}$, where U_{coul} is the Coulomb potential.

2.3.4.2 Particle Mesh Methods

The Ewald summation is a great way of implementing the long-range forces, but the drawback is that the method scales as $\mathcal{O}\left(N^2\right)$, or $\mathcal{O}\left(N^{3/2}\right)$ at best [50]. The most time consuming part is the Fourier transformation in 2.64, therefore several methods were developed where the original problem is rewritten, such that the fast-Fourier-transformation (FFT) can be used, meaning that the computational effort of the Fourier part is reduced to $\mathcal{O}\left(N \log N\right)$. However, the FFT is a grid transformation, which therefore means that the charges need to be discretised on a mesh, and the resulting error must be minimised/controlled. This type of methods are called Particle Mesh methods, and although several of them exist, they follow similar ideas but are different in detail. The first particle mesh method was developed by Hockney and Eastwood [89] and is called the *Particle-Particle/Particle-Mesh (PPPM)* method. Variants of this method are the *Particle Mesh Ewald (PME)* method [47], and the *Smooth Particle Mesh Ewald (SPME)* method [58]. However, in two papers by Deserno and Holm [50, 51] it was shown that the PPPM method is the most accurate one, and that it has a way of estimating the error. Also, the new ideas used in the variants of the method can as well be applied to PPPM, so this method remains the preferred method to be used [50].

Next, the basic steps involved in the PPPM method are briefly explained; for a more thorough introduction the reader is referred to the references mentioned before. Similar to the Ewald summation method, the general idea of the PPPM method is to split the interparticle force into an exponentially decaying short-range force, and a smoothly varying long-range part that should be equal to the actual value at large distance and that can be determined efficiently with an FFT. The short range part can be computed with a direct particle-particle (PP) force summation, where a certain cutoff radius can be used and can be evaluated using the cell list method, as discussed in Section 2.3.3. The long-range part, on the other hand, is determined on a particle mesh (PM), meaning that the system first needs to be mapped onto a mesh. Once discretised, the charge-potential can be computed using the FFT. After this is done, the appropriate values of a given property, e.g. the electric field, need to be determined using the mesh. The final step is to interpolate the (long-range) grid data back to the atomic positions. Each step incurs a numerical error. However, by a clever optimisation of the method, the inaccuracies in the individual steps can be made (almost) self cancelling overall [89]. The four steps are summarised in Figure 2.15.

The first step, i.e. the mapping of the charges on a grid representation of charge density at discrete locations, basically means that the charge of the atom at position $\vec{r}_i$ is assigned to several mesh points $\vec{x}_p$, per coordinate direction. This is accomplished by means of a charge assignment function $W\left(\vec{x}\right)$ that assigns the fraction of charge to the mesh point $\vec{x}_p$ due to

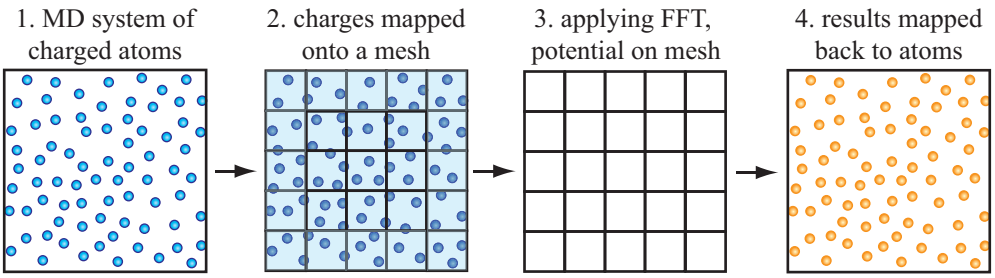

Figure 2.15: The four steps of obtaining the long-range part of the force using a Particle-Mesh method. 1. the MD system with charged atoms. 2. The charges are mapped onto a mesh. 3. the charge-potential on the mesh is computed using the FFT. 4. The charge-potential is mapped back to the atom positions.

a (unit) charge at position $\vec{r}_i$. The mesh-based charge density can then be written as [89]:

$$\rho_M\left(\vec{x}_p\right) = \frac{1}{h^3}\sum_{i=1}^{N} q_i W\left(\vec{x}_p - \vec{r}_i\right) \tag{2.67}$$

where $h = {}^{L}/_{N_M}$ is the grid spacing where N_M are the number of mesh points in each direction. In the original PPPM method [89] the charge assignment function was chosen to be a function that distributes the charge of a single atom between its P nearest mesh points (i.e. function is said to be of the P^{th}-order). The Fourier transformed charge assignment function in that case can be written as [89]:

$$\tilde{W}^{(P)}\left(\vec{k}\right) = h^3\left(\frac{\sin\left(\frac{1}{2}k_x h\right)}{\frac{1}{2}k_x h}\frac{\sin\left(\frac{1}{2}k_y h\right)}{\frac{1}{2}k_y h}\frac{\sin\left(\frac{1}{2}k_z h\right)}{\frac{1}{2}k_z h}\right)^P \tag{2.68}$$

The first part of the PPPM method is straightforward, which cannot be said about the next step which is solving for the potential field and obtaining the derivatives of the potential field to obtain the electric field and corresponding forces on the mesh. In the original Ewald summation the Fourier part of the summation is given by 2.64, but this equation now needs to be rewritten, so it can be solved on a mesh with an FFT. Each of the three particle mesh methods mentioned above uses a different technique to achieve this. The easiest method is to substitute the transformed charge density with a FFT of the discretised charge density, which is done in the PME method [47]. Although this is easy to implement, it is not necessarily the most accurate solution possible [50]. The PPPM method uses a clever technique applied to the Fourier transform in which case the difference between the results of the computation on the mesh (with the mapped charges) and the continuum problem (Poisson's equation with the real positions of the charges) is minimised. This is accomplished by multiplying the Fourier transform of the mesh based charge density with the Fourier transform of the so-called *optimised influence function* $\vec{G}_{opt}$ which is noted as $\hat{G}_{opt}$ and equals [89]:

$$\hat{G}_{opt}\left(\vec{k}\right) = \frac{\tilde{D}\left(\vec{k}\right)\cdot\sum_{\vec{m}\epsilon\mathbb{Z}^3}\tilde{U}^2\left(\vec{k}+\frac{2\pi}{h}\vec{m}\right)\tilde{R}\left(\vec{k}+\frac{2\pi}{h}\vec{m}\right)}{\left|\tilde{D}\left(\vec{k}\right)\right|^2\left[\sum_{\vec{m}\epsilon\mathbb{Z}^3}\tilde{U}^2\left(\vec{k}+\frac{2\pi}{h}\vec{m}\right)\right]} \tag{2.69}$$

where:

$$\tilde{U}\left(\vec{k}\right) = \frac{\tilde{W}\left(\vec{k}\right)}{h^3} \tag{2.70}$$

and $\tilde{D}\left(\vec{k}\right)$ is the Fourier transform of the employed differentiation operator that is discussed next. The term $\tilde{R}(k)$ in 2.69 is the Fourier transform of the long-range part of the

interparticle force, and in the case of Coulombic interaction it is defined as [89]:

$$\tilde{R}\left(\vec{k}\right) = -\imath\vec{k}\frac{4\pi}{k^2}\exp\left(-\frac{k^2}{4\alpha^2}\right) \tag{2.71}$$

However, this term can be changed when different long-range forces need to be computed, as will be shown later. Please note that the influence function in this form does not depend on the positions of the atoms and therefore it only needs to be computed for each FFT mesh. Once the electric potential has been solved, the electric field, defined as $\vec{E} = -\nabla\phi$, can be computed by taking the finite differences on the mesh. However, another way to differentiate the potential is to do it already in Fourier space. This can be accomplished by multiplication with (minus) the Fourier transform of the employed differentiation operator, $\imath\vec{k}$. In that case the mesh based electric field, $\vec{E}_M$, is computed by [50]:

$$\vec{E}_M\left(\vec{x}_p\right) = \overleftarrow{\text{FFT}}\left[-\imath\vec{k}\times\overrightarrow{\text{FFT}}\left[\rho_M\right]\times\hat{G}_{opt}\right]\left(\vec{x}_p\right) \tag{2.72}$$

where the arrows indicate whether it is the backward or forward FFT that is computed, and consequently $\tilde{D}\left(\vec{k}\right) = \imath\vec{k}$ in the optimised influence function. After this operation the electric field is known on the mesh.

The final stage of the particle mesh method is to interpolate the discrete values on the mesh back to the actual positions of the atoms and to compute the force on each atom. This can be accomplished in a similar manner as what was done when assigning the charges to the mesh, e.g. by some assignment function, as follows [89]:

$$\vec{f_i} = q_i\sum_{\vec{x}_P\in\mathbb{M}}\vec{E}_M\left(\vec{x}_P\right)W\left(\vec{r}_i-\vec{x}_P\right) \tag{2.73}$$

where the summation is done over all mesh points in the mesh $\mathbb{M}$. For practical reasons the assignment function in this equation is chosen exactly the same as the one used before, however it can be proven [89] that this is also one of the minimum requirements to obtain conservation of momentum.

This concludes the four steps involved in the computation of the long-range part in the PPPM method. By adding the short-range term and the long-range term, the total force on each atom is known. However, the value of the force is only accurate enough, i.e.. with errors less than 10^{-5}, if the proper values are used for the model parameters like α and FFT mesh size. A detailed discussion of this is given by Deserno and Holm [51].

So far, mainly Coulombic interaction, $(1/r)$, is discussed, but it is also possible to use the PPPM method for other interactions of the type $(1/r^p)$ by only changing small parts of the functions used in the PPPM method [175]. The force splitting in the PPPM method is based on the following trivial identity:

$$\frac{1}{r^n} = \frac{f\left(r\right)}{r^n} + \frac{1-f\left(r\right)}{r^n} \tag{2.74}$$

where the first part is the exponentially decaying short-range term and the second part the slow varying long-range term. These requirements leave many choices for the function $f\left(r\right)$. Essmann et al. [58] propose such a function. The non-trivial derivation of this function is given in Appendix B. The result is a general expression for $1/r^p$:

$$\frac{1}{r^p} = \pi^{3/2}\beta^{(p-3)}\int_{\mathbb{R}^3} f_p\left(\frac{\pi\vec{u}}{\beta}\right)\exp\left(-2\pi\imath\vec{u}\cdot\vec{r}\right)d^3\vec{u} + \frac{g_p\left(\beta r\right)}{r^p} \tag{2.75}$$

where

$$f_p\left(x\right) = \frac{2\left(x\right)^{p-3}}{\Gamma\left(p/2\right)}\int_x^\infty s^{2-p}\exp\left(-s^2\right)ds \tag{2.76}$$

and

$$g_p\left(x\right) = \frac{2}{\Gamma\left(p/2\right)}\int_x^\infty s^{p-1}\exp\left(-s^2\right)ds \tag{2.77}$$

In these equations, Γ is the (Euler) Gamma function, defined as:

$$\Gamma\left(z\right) = \int_0^\infty t^{z-1}\exp\left(-t\right)dt \tag{2.78}$$

Table 2.1: The functions f_p and g_p for several values of p. Here $\Gamma[a,z] = \int_z^\infty t^{a-1}\exp(-t)\,dt$

p	$f_p(x)$	$g_p(x)$
1	$\exp(-x^2)/x^2\sqrt{\pi}$	$\mathrm{erfc}(x)$
2	$\frac{\sqrt{\pi}}{x}\mathrm{erfc}(x)$	$\Gamma[0,x^2]$
3	$2\Gamma[0,x^2]/\sqrt{\pi}$	$\mathrm{erfc}(x) + 2x\exp(-x^2)/\sqrt{\pi}$
6	$\frac{1}{3}\left[\left(1-2x^2\right)\exp\left(-x^2\right) + 2x^3\sqrt{\pi}\mathrm{erfc}(x)\right]$	$\exp\left(-x^2\right)\left(1+x^2+\frac{1}{2}x^4\right)$

These equations give the long-range part that can be solved with the particle mesh method and the short-range part that can be solved with a particle-particle method for any $1/r^p$ type of interaction. The only requirement is to solve the f_p and g_p functions, which in most cases can be done analytically, and compute the Fourier transformed long-range part of the interparticle force, $\tilde{R}(k)$. Table 2.1 gives the function f_p and g_p for several values of p.

For the Coulombic interaction it can easily be seen how the comparable values of f_p and g_p appear in the Ewald summation 2.63, 2.64. However, this method can now also be used to compute the long-range Van der Waals attraction term in the Lennard-Jones potential when the functions of $p = 6$ are used. The strength of this method is that there is no need for using a cutoff, where the cutoff can cause significant errors in system properties, like surface tension. The $1/r^6$ term is split in the long-range and short-range part according to B.10, where the long-range part of the potential and the resulting part of the force by differentiation, is solved with the particle mesh method explained in this section. The short-range part of the potential is $g_6(\beta r)/r^6$. The force, including the $1/r^{12}$ term, can be computed by $\vec{f} = -\nabla U$:

$$\vec{f}_{ij,SR} = \frac{48\epsilon}{\sigma^2}\left(\frac{\sigma}{r_{ij}}\right)^2\left[\left(\frac{\sigma}{r_{ij}}\right)^{12} - \frac{1}{2}g_6(\beta r)\left(\frac{\sigma}{r_{ij}}\right)^6 + \frac{1}{12}r_{ij}\frac{dg_6(\beta r)}{dr}\left(\frac{\sigma}{r_{ij}}\right)^6\right]\vec{r}_{ij} \qquad (2.79)$$

and is easily incorporated in the normal force loop. In section 2.4.4 several droplet formation simulations are done to compare the influence of the cut-off value on the droplet break-up and formation. These results will be compared to a simulation where the long-range techniques are also used for the Van der Waals attraction term while using a small cut-off radius.

2.3.4.3 Long-range Forces in Non-Periodic Domains

The methods to compute the long-range forces discussed so far assume an infinite periodic domain in *all* directions. However, in certain situations, for example when there are atomistic boundaries present, the domain is only periodic in two dimensions or even only in one dimension. In these cases the Ewald sum or Particle Mesh methods do not compute the appropriate solution and therefore need to be adapted.

One of the first methods to compute the Coulombic interaction in a domain with two periodic dimensions, was developed by Parry [144, 145], Heyes et al. [88], and de Leeuw et al. [119]. This resulted in a 2D Ewald method, where the reciprocal space is defined in two directions. However, the disadvantage of this method is that no error estimate is available, and the method is very impractical for larger systems (i.e. $N > 10^3$). The next development was by Spohr [179], who used the conventional 3D Ewald technique with a different boundary condition and added an empty space in the direction that is not periodic with the aim to minimise the (artificial) interaction between the replicated domains in the reciprocal space. The advantage of this method is that error estimates are available. However, it was shown by Yeh and Berkowitz [213] that the error is only reduced slowly when increasing the size of the empty space. Also, depending on the splitting parameters, at some point the increase of the empty space no longer result in a smaller error. To improve the error convergence, a correction term, consisting of the total dipole moment of the original MD system was introduced [213]. Their method showed that the same accuracy can be achieved with a much smaller empty space, but the error does not vanish and reaches a plateau [30]. Improved accuracy is obtained by the method of Arnold et al. [11, 99], where an additional correction, the so-called electrostatic layer correction (ELC), is applied that removes the unwanted

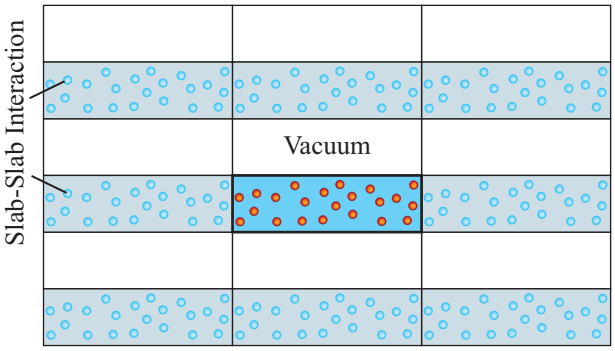

Figure 2.16: The idea behind the 2D slab ELC method [11]. A vacuum is added to the (periodic) MD system and resulting slab-slab interaction is corrected.

interactions of the system with the replicated layers along the z-axis. This method is also easily incorporated into the PPPM method [11] and only requires slightly more computation time than the conventional method where all three dimensions are periodic. Figure 2.16 illustrates this method.

Similar methods were developed for the case where only one dimension is periodic. using similar techniques that are used to develop the 2D Ewald method, Porto [153] developed a 1D Ewald method. However, this method has a $\mathcal{O}\left(N^2\right)$ scaling, and again no a-priori error estimate is available. On the other hand, there exist methods [77, 79] that are based on the Lekner method [120]. Using the Lekner method the summation of the Coulomb interactions is done entirely in reciprocal space and is faster than the 1D Ewald method, but it becomes numerically unstable for small distances between the charges. Arnold et al. [10] developed a method, called MMM1D, that uses the Lekner approach for charge interaction at large distances, while for charges at close distances a rapidly convergent series is used. In this way the numerical instability is avoided.

2.3.5 Modelling Argon and Reduced Units

In the previous sections, the general theory of molecular dynamics and some advanced techniques were discussed. In this section it will be explained how the method can be used to obtain results using a certain substance or material. One of the earliest successful molecular dynamics simulations were done in 1964 by Rahman [155], where the properties of liquid argon were investigated. Together with the fact that the forces between argon can be simulated very accurately by a (simple) Lennard-Jones potential, argon remains the liquid of choice when the molecular dynamics method is discussed in the literature. However, the Lennard-Jones potential also gives good results for other liquids, where the interaction between the atoms is well-presented by long-range van der Waals attraction and short-range Pauli repulsion. However, the Lennard-Jones 12-6 potential has only two parameters, ϵ and σ, which need to be determined by comparing simulations and, for example, the values of the pressure obtained from empirical equations of state, or real properties that can be measured in experiments. These parameters are different for every liquid that is modelled, and therefore the equations derived so far also change accordingly. For example, Table 2.2 gives some commonly used values for ϵ and σ for different noble gasses [20]. Please note that the value for ϵ is given in units of temperature by using the Boltzmann constant, $k_B = 1.3807 \times 10^{-23} JK^{-1}$.

However, there is a way to rewrite (or scale) the equations in such a way that they are identical for every liquid, i.e. the molecular dynamics simulation code can be made independent of the actual values of σ and ϵ. This is done by using suitable dimensionless units, hereafter called *reduced units*. They are based on the choice of σ, m, and ϵ as the units of length, mass, and energy, respectively. This is shown in Table 2.3 where the formula of how different physical units can be converted into reduced units is given. By using these reduced units, the Lennard-Jones 12-6 potential, 2.50 can be rewritten as:

Table 2.2: The Lennard-Jones parameters ϵ and σ, which can be used to accurately simulate a certain substance

Element	ϵ/k_B (K)	σ (nm)
Ne	35.7	0.2789
Ar	122.4	0.3432
Kr	170.0	0.3675
Xe	234.7	0.4009

Table 2.3: The conversion of several sampled variables between reduced units and physical units as a function of the Lennard-Jones potential parameters σ and ϵ. The reduced units are indicated with an asterisk. The right column shows how the *one* reduced units corresponds to physical units for argon, given the mass of an argon atom, $m = 39.948 \times 1.6605 \times 10^{-24} g$.

Variable	Formula	for argon (1 reduced units =)
Time	$t \rightarrow t^* \sqrt{m\sigma^2/\epsilon}$	2.1502×10^{-12} s
Density	$\rho \rightarrow \rho^* m/\sigma^3$	1.6410×10^3 kg/m^3
Velocity	$v \rightarrow v^* \sqrt{\epsilon/m}$	159.61 m/s
Temperature	$T \rightarrow T^* \epsilon/k_B$	122.4 K
Force	$f \rightarrow f^* \epsilon/\sigma$	4.9240×10^{-12} N
Stress	$\sigma \rightarrow \sigma^* \epsilon/\sigma^3$	41.805 MPa
Viscosity	$\mu \rightarrow \mu^* \sqrt{\epsilon m}/\sigma^2$	0.8989×10^{-4} Pas

$$V^* = 4\left(\left(r_{ij}^*\right)^{-12} - \left(r_{ij}^*\right)^{-6}\right) \tag{2.80}$$

while the corresponding force, 2.51, can be rewritten in reduced units as:

$$\vec{f}_{ij}^* = 48\left(r_{ij}^*\right)^{-2}\left[\left(r_{ij}^*\right)^{-12} - \frac{1}{2}\left(r_{ij}^*\right)^{-6}\right]\vec{r}_{ij}^* \tag{2.81}$$

The parameters with an asterisk are the reduced values of the corresponding parameter. After a certain simulation with any liquid is performed, the real (physical) values of, for example velocity, pressure, or temperature, can be computed using the appropriate conversion factors. For argon these are provided in the third column of Table 2.3. Please note that where specified otherwise, all results of the simulations are given in reduced units.

In Section 2.3.1 the Verlet algorithm was introduced, which is used to integrate the MD equation in time. However, so far nothing was mentioned how large the time step should be. Typical time steps for the Verlet integration are taken as $\Delta t^* = 0.005$, which corresponds to $\Delta t = 1.0751 \times 10^{-14} s$ for argon. This incredibly small time step means that for 1 nanosecond in physical time, almost 10^5 iterations (and interparticle force computations) need to be performed. Luckily, this small time step does not limit the fact that molecular dynamics is still a very useful technique in obtaining macroscopic (long time averaged) properties.

This concludes this section where reduced units are introduced and what the meaning of the Lennard-Jones parameters are. In the next section, it will be shown how more realistic fluids, like water, can be modelled. These models still use the Lennard-Jones potential and its parameters, but also the electric forces between atoms and molecules is included.

2.3.6 Modelling Water

So far the modelling of atoms and molecules was limited to non-polar cases, especially noble substances like argon, where the Lennard-Jones potential can be applied reliably. Although numerous experiments were done with liquid argon, a substance like water does have many practical technical applications. Therefore, to study phenomena encountered in these applications, a model of water for use in a molecular dynamics code would be desirable. On the other hand, water does have some unique properties, and appropriate models of

Table 2.4: The model parameters for the SPC [17], SPC/E [16], TIP3P [101], TIP4P [102], TIP4P/Ew [92], TIP4P/2005 [2], and TIP5P [125] water model.

Water model	SPC	SPC/E	TIP3P	TIP4P	TIP4P/Ew	TIP4P/2005	TIP5P
Type	3-sites	3-sites	3-sites	4-sites	4-sites	4-sites	5-sites
ϵ_{OO} (kJ/mol)	0.650	0.650	0.6364	0.6480	0.680946	0.7749	0.6694
σ_{OO} (Å)	3.166	3.166	3.15061	3.15365	3.16435	3.1589	3.12
q_1 (e)	+0.410	+0.4238	+0.4170	+0.5200	+0.52422	+0.5564	+0.2410
q_2 (e)	−0.820	−0.8476	−0.8340	−1.0400	−1.04844	−1.1128	−0.2410
L_{OH} (Å)	1.0000	1.0000	0.9572	0.9572	0.9572	0.9572	0.9572
L_{OD} (Å)	-	-	-	0.15	0.125	0.1546	0.70
θ (^{o})	109.47	109.47	104.52	104.52	104.52	104.52	104.52
φ (^{o})	-	-	-	52.26	52.26	52.26	109.47

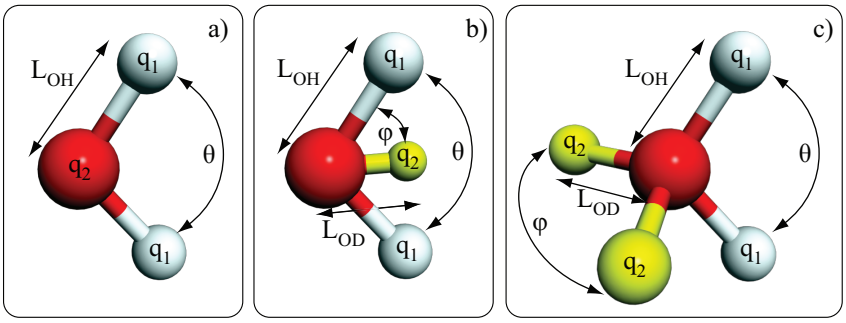

Figure 2.17: The water models that use three (a), four (b), and five (c) interaction sites.

water could give more insight into these properties. However, because of the complexity of the behaviour of water and the accompanied difficulty of modelling this together with the required computational time, some trade-offs are required. Roughly speaking, the different models in use today can be distinguished by how many sites (points of charge or atoms) are included in the model and the inclusion of flexible bonds and/or polarisation. The models for water introduced next all have a number of things in common. Firstly, the interaction due to charge is computed with Coulombic interaction, while the normal interatomic interaction is modelled with the (normal) Lennard-Jones potential. Secondly, the length of the bonds between the oxygen atom and the hydrogen atoms is kept constant, which is also applied to the angle between the atom bonds. However, the models differ in where the interactions takes place. Figure 2.17 gives the general impression of the different water models, while table 2.4 gives the parameters for several models explained next.

The simplest model has *three* sites, where the interaction sites correspond to the relative positions of the oxygen atom and the hydrogen atoms. All atoms have a charge, but only the oxygen atom has Lennard-Jones parameters. The most commonly used 3-site models include the *SPC* [17] (Simple Point Charge) model, the *SPC/E* [16] (Extended version of the latter), and the *TIP3P* [101] (Transferable Intermolecular Potential with 3 Points). The difference between these models are the Lennard-Jones parameters that are used, the bond length (O-H), and the HOH-angle. Consequently, the models also predict different bulk properties and have different phase diagrams. Therefore, the model in the MD code must be chosen carefully.

The models that use *four* interaction sites introduce a dummy atom that carries (and therefore moves) the charge of the O-atom away from the O-atom itself. In doing so, the electrostatic distribution around the water molecule can be modelled more accurately. Possibly the oldest realistic model of water is a 4-site model and was developed in 1933 by Bernal and Fowler [19]. Because this model (the BF-model) was developed before the onset of computer simulations, the parameters of the model are not optimised, and consequently the bulk properties of water are not simulated correctly. However, nowadays the model parameters

can be optimised to match a range of bulk properties, especially density, radial distribution functions, and diffusion constant, as best as possible. Popular 4-site models include the *TIP4P* [102] model and several variants that specialise in a certain property to match. For example, the *TIP4P/2005* model [2], where the focus is on the phase diagram of water, or the *TIP4P/Ew* model [92], which is optimised to work with Ewald summations. Again the models differ only in their parameters, which now also includes the length of the bond between the dummy atom and the oxygen atom and the angle at which the dummy atom is placed with respect to the H-atom.

There are also several models that use more than four interaction sites. However, with every added site the computational effort increases, so they are only worthwhile if they actually improve the accuracy of the target property that is modelled. The idea of the 5-site water model, for example the *TIP5P* model [125], is to add two charged dummy atoms that again move the charge away from the O-atom. These negative charges on the outside try to model the Lone pair of the water molecule [125]. Because the water molecule itself is electrically neutral, each dummy atom has the same value, but a different sign of the charge as the H-atoms. And finally the 6-site model [135] includes the dummy atoms of both the 4-site and 5-site models.

It was mentioned in the introduction that all models discussed here put constraints on the length of the bonds and/or angles between the atoms. There are a couple of advantages in doing so. The main obvious advantage is the maximum allowed time step for the accurate integration of the equations of motion. The fastest motions in the water molecule are the bond vibrations. Taking away this motion, by fixing the bond-length, allows an increase in the time step. Connected to this, the (actual) vibrating frequency can be so high that the classical approximation (Newtonian) is no longer valid anymore, and quantum effects become important. So, to avoid this to happen, holonomic (fixed) constraints are used. The most popular algorithm used is called *SHAKE* [165]. Generally speaking, in MD the constraint algorithm computes the constraint forces on the atoms in such a way that the length of the bond remains exactly fixed after the end of each time step used in the MD sode. This (non-linear) computation is done iteratively, and therefore needs the specification of a tolerance or maximum number of iterations. SHAKE uses the Verlet integration scheme, although other versions of SHAKE, like *RATTLE* [9] or *WIGGLE* [118], make use of the velocity Verlet method. In the case of the water molecule, where (only) two or three bonds need to be fixed, analytical versions of the constraint algorithms are available. This algorithm is called *SETTLE* [131]. By using these constraint algortihms, the commonly accepted maximum allowed time step can be taken between $\Delta t = 1 - 2\,\mathrm{fs}$ in most MD simulations of water, which is about five to ten times smaller than in typical MD simulations of argon. But unforunately, if larger time steps in water simulations ($t > 4\,\mathrm{fs}$) are used, this will lead to severe energy drifts and incorrect statistical properties [97], because of instabilities in the integration of the equations of motion. However, there are ways, called multiple time stepping intergrators, where an *approximate* MD simulation can be performed with larger time steps [97, 124]. These methods work on the principle that the forces are split into fast and slow components, where the fast component is evaluated more frequently that the slow component. Using this method allows for time steps of $4 - 6$ fs, or if a (targeted) Langevin coupling [124] is employed, time steps of up to 16 fs are possible.

In the previous section the reduced units were explained and how they are applied when simulating argon. Although the same can be applied for the simulation of water, it is more common to work in physical, but specially scaled units. For example, the unit of time is femtoseconds, (fs), while length is expressed in Angstroms, (Å), and mass in atomic weights, or $(\mathrm{g/mol})$. A frequently used cutoff radius for the interatomic interactions is taken as $r_{cut} = 1\,\mathrm{nm}$, which is slightly more than the cutoff value used for argon (see Section 2.3.3).

As will become clear in Chapter 3, the material property that is important for the coupling is mainly the viscosity. However, most of the aforementioned water models were never optimised for this property, and therefore they give wrong values of the viscosity at a certain temperature compared to experimental values. In Section 2.4.5, the viscosity of several water models is given as a function of temperature. From these results, the best fitting water model is selected for the MD/continuum coupling.

2.4 Results

2.4.1 CFD: Electrokinetic Flow Inside a Nanofluidic Device

In this section * it is shown how the pure continuum equations of motion for electrokinetic flow can be used to successfully simulate nanofluidics, while the underlying phenomenon is basically on the molecular scale.

In the last few years, a large number of fluidic devices are being studied, driven by the motivation to develop bio-medical applications that consume small amounts of reagents and analytes. Among these, one important explored direction is the development of devices that allows for single molecule detection (SMD) and single molecule manipulation (SMM) [62, 180]. If a reliable SMD method exists, it will allow for a sensitive detection of very low concentrations of the tested fluid, a key issue in medical diagnostics. The main challenge in SMD is to optically detect small volumes. Different methods have been described to detect small volumes, among others is to use submicrometer-sized fluidic channel [180]. On the other hand, the major challenge in SMM is the precise manipulation of single (large) molecules, which requires a sufficient control of the molecule and the flow. Many lab-on-a-chip applications use electroosmotic flow to solve this challenge. In order to perform an efficient study on single molecules, both the challenges on SMD and SMM have to be solved simultaneously. On the other hand, nanofluidic channels and Lab-on-a-Chip devices have also been used and developed to perform investigation on DNA molecules, an interesting issue in genomic studies and polymer science [188]. The information obtained from the DNA sequence is particularly important in genomic applications, as was highlighted in the Human Genome Project [198]. A crucial step in the sequencing process is the size-based separation of DNA fragments under an electric field, which conventionally requires a sieving tool due to the size-independent electrophoretic DNA mobility in free-solution [202]. In practice, the sieving tool can be in the form of a drag-tag [129], gel matrices [202], dielectrophoretic traps [148, 157], membranes [183], or microfabricated periodic structures [13, 31, 37, 67, 66, 83, 95, 94]. The idea of a DNA separation in Lab-on-a-Chip that does not require any sieving tool is very appealing, because extra preparatory steps (as in the cases of drag-tag and gel matrices) or complicated sieve designs (as in the cases of dielectrophoretic traps, membranes, and periodic microstructures) could then be avoided. On top of that, the DNA separation scheme should be able to run continuously, which would allow high throughput [56]. Some interesting observations have been reported on electrokinetic phenomena of DNA in nanofluidic channels [15, 111, 161], and more recently sieve-less DNA separation in nanochannels have been reported, although they are operated in a non-continuous manner [41, 146]. In this section, the electrokinetic pathlines of DNA molecules in U-turn nanofluidic channels are studied, in order to develop a *continuous*, sieve-less, size-based DNA separation.

The U-turn nanofluidic device that is used and modelled in this section is shown in Figure 2.18. The device has a 100 micrometre wide straight inlet channel, which is connected to a semi-circular chamber with a radius of 1 millimetre. This chamber is connected to twelve 100 micrometre wide sub-outlets collecting the fluid out of the chamber into a 273 µm wide main outlet. However, the height of the device is only 150 nanometres. This small depth of the device allows to optically detect fluorescence from a volume as small as 270 attoliters (al), which is adequate for sensitive SMD. The molecule-sorting scheme relies on electrophoresis and dielectrophoresis, combined with the branched U-turn shape of the device, while the transport is accomplished with electroosmosis. The externally applied electric field is accomplished by electrodes at the inlet and the outlet of the device, which generates the electric potential (voltage) difference. Here the numerical details will be highlighted, therefore for detailed information regarding the fabrication of the device and the used experimental setup, the reader is referred to [141], [142] for the electroosmotic and electrophoretic setup, and [143] for the di-electrophoretic setup.

Before dielectrophoretic behaviour of DNA inside the device is investigated, first the electroosmotic flow and the electrophoretic behaviour of simple solid particles (i.e. beads) are examined. The device has a depth of only $h = 150$ nanometres, whereas the working fluid is distilled water, for which the Debye length is of the order of 1 micrometre. This means that $\kappa h \sim 0.1$, where κ is defined in 2.7. Furthermore, the Reynolds number of the flow is of the order of 10^{-3}. These conditions meet the requirements needed to invoke the thick EDL

*Published as: [126, 142, 143]

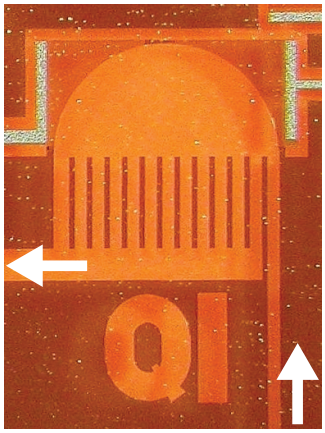

Figure 2.18: Microphotograph of the nanofluidic device which is investigated. Electrodes are used to create electroosmotic, electrophoretic, and dielectrophoretic flow inside the device to manipulate and sort DNA molecules.

approach, and therefore this device can be used to verify the approach. To visualise the electroosmotic flow inside the device, a solution of 110 nanometres diameter fluorescent beads are used as the tracer-particles. These beads are set in motion due to electrophoresis, so the visualisation is the sum of electroosmotic and electrophoretic flow. In Section 2.1.2.1 it was established that the electroosmotic velocity profile in the thick EDL case has a parabolic shape in the direction of the height of the device, but an external observer is more likely to only see the beads moving. Therefore a connection between the velocity profile and the velocity of the beads is needed. One way of doing this is to use the vertically-averaged velocity inside the channel. As the diameter of the beads is of the order of the channel height, these beads do not flow at this averaged velocity, given to their limited size, but travel at a certain velocity between the average and the maximum velocity. For a parabolic velocity profile the maximum velocity is $\frac{3}{2}$ of the mean velocity, therefore 2.10 is modified with a factor, in which $1 \leq C_1 \leq \frac{3}{2}$, and that is dependent of the size of the bead. Meanwhile, the electrophoretic velocity of the beads, which have a non-zero zeta potential ζ_p, can be computed using 2.11 in the thick EDL case. This means that in total, the observed velocity of the bead in the experiments will be:

$$\vec{u}_{observed} = -\frac{\epsilon\epsilon_0}{\mu}\vec{E}\left(\frac{(\kappa h)^2}{12}C_1\zeta_w - \frac{2}{3}\zeta_p\right) \tag{2.82}$$

However, because the individual parameters are not known exactly it is more convenient to rewrite this equation into:

$$\vec{u}_{observed} = -\frac{\epsilon\epsilon_0\zeta_{eff}}{\mu}\vec{E} \tag{2.83}$$

where ζ_{eff} is an effective zeta potential that incorporates the effect of the wall zeta potential, the finite size of the beads, and the zeta potential between the beads and the liquid. The local velocity is still directly related to the local electric field inside the channel, and it has the same format as the Helmholtz-Smoluchowski relation 2.8, where the zeta potential is substituted by an effective zeta potential. Because of the dimensions of the device, the electric field is simply computed using Laplace's equation 2.9 for the electric potential, and is solved using the finite element method discussed in Section 2.2.2. The boundary conditions for this equation depend on the voltage applied to the electrodes. However, note that, because of the governing equations, only the strength of the electric field depends on the value of the applied voltage, while the shape of the electric field is invariant. Consequently, this also applies to the velocity vectors. To make certain that the numerical geometry matches that of the physical geometry, a photo of the structure taken through a microscope is used as reference. Further, this photo does not show any visible rhougness at the walls

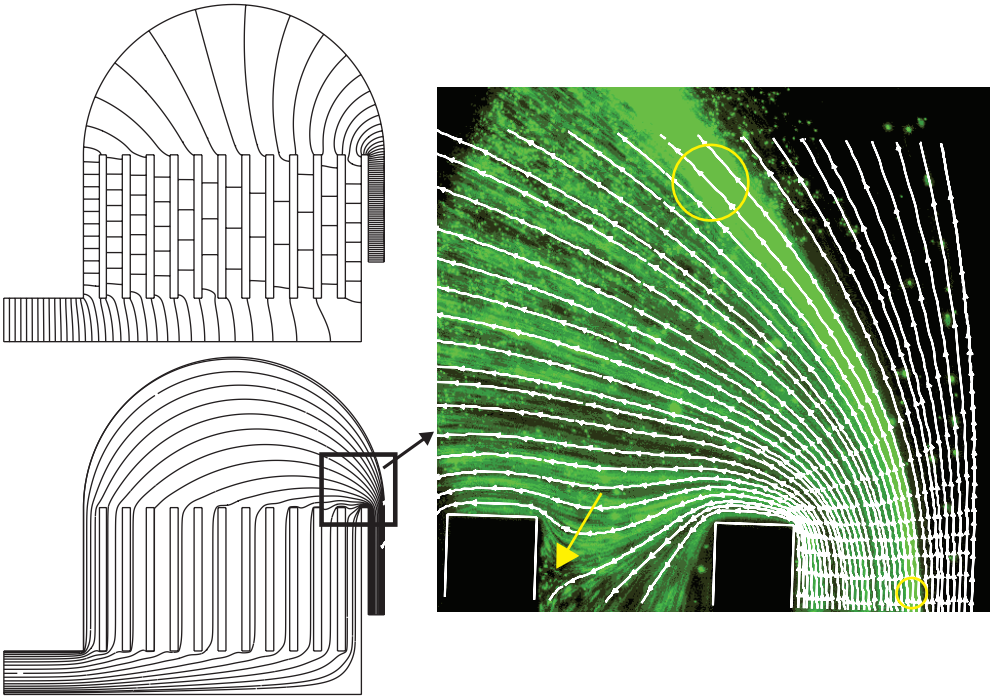

Figure 2.19: The numerical results of the electric field lines (top left), the resulting pathlines (bottom left), and the details near the entrance where the pathlines of the beads obtained from the measurement with an applied voltage potential of 300 Volts are compared with the numerical results represented by the continuous pathlines (right)

of the device. Therefore, to aid the numerical model, the walls are assumed to be perfectly smooth.

In [142] an electroosmotic flow analysis of the nanofluidic device was presented for a voltage potential of 100 Volts applied between the electrodes. The images of the tracer-particles in the fluid were captured, and the streamlines and the velocity distribution of the fluid in the device were analyzed and computed. This provided the opportunity to determine how well the actual flow field from the measured tracer motion can be determined. Also, when all the images captured during the experiment are added up, they create a single image of the (real) pathlines of the beads. In general, the measured and computed pathlines agreed with each other. The observed deviations are most likely the result of Brownian motion, which is not included into the numerical simulation. To verify this, additional experiments with voltage potentials of 30 V and 50 V between the electrodes were performed. The results showed that the deviations were indeed caused by Brownian motion. Higher potential imply a higher velocity of the tracer particles, and when the velocity is significantly higher than the rms drift velocity of the Brownian motion, the drift of the particles from their respective pathlines becomes negligible. Therefore, the meaurements are repeated with a electrode potential difference of 300 V. The new results of this meaurement are shown in Figure 2.19. Here the top-left part shows the electric fieldlines inside the device, the bottom-left part shows the resulting theoretical pathlines, while the right figure shows the details of the comparison between the numerical and the experimental pathlines near the entrance. In this case the measured streamlines do coincide with the pathlines computed with the thick EDL approach. This can especially be observed for the several diverging streaks that start at the inlet at the right (indicated by the circle in the figure). The pathlines diverge in similar ways. The experimental pathlines in the first sub-outlet show a part that goes into the sub-outlet and a part that escapes from the sub-outlet with a stagnation point indicated by an arrow in Figure 2.19. This is also seen in the computed pathlines. Further, the computed pathlines just above the walls between the inlet and the sub-outlets follow a similar path as

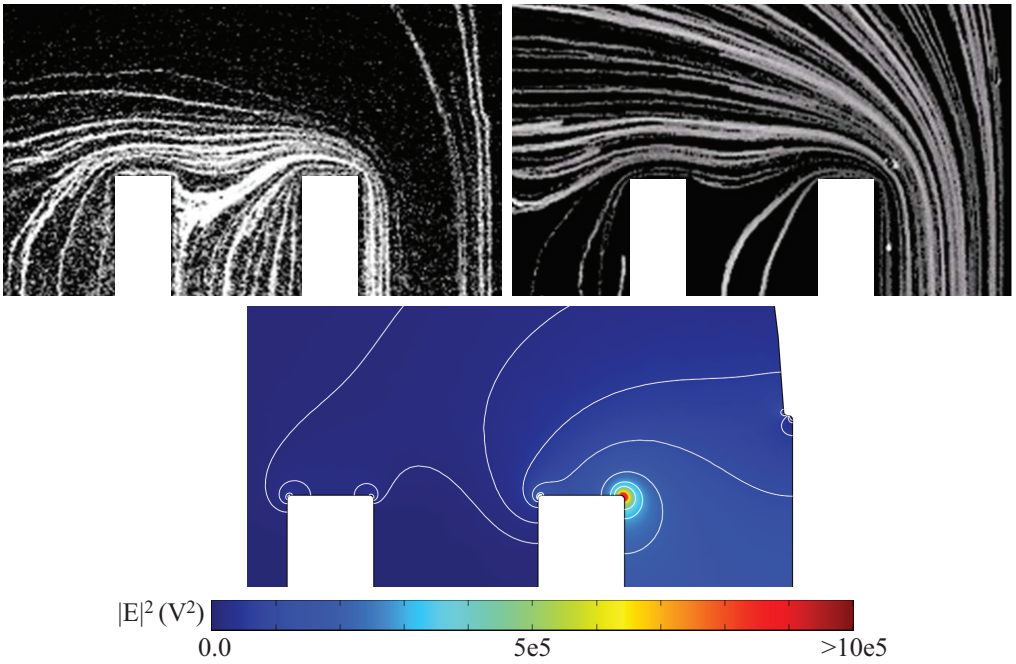

Figure 2.20: Two examples of the sets of pathlines of DNA molecules, where the value of the electric field is $V = 15.3$ V (DC) at the top, and the numerical results of electric field intensity $\left|\vec{E}\right|^2$ at the bottom. The top left figure shows the pathlines of λ-DNA molecules, while the top right figure shows the pathlines of T4GT7-DNA molecules. For the numerical results several iso-level contours of $\left|\vec{E}\right|^2$ are also shown in the figure ($\left|\vec{E}\right|^2$ equal to $0.1 \cdot 10^5$, $0.2 \cdot 10^5$, $0.5 \cdot 10^5$, $1 \cdot 10^5$, $2 \cdot 10^5$, $3 \cdot 10^5$, $4 \cdot 10^5$, $5 \cdot 10^5$, and $10 \cdot 10^5$) for visual aid.

the experimental pathlines. Therefore it can be concluded that the simple equations of the thick EDL approach do compute the correct type of flow behaviour inside this nano-sized device.

The next set of experiments are done with DNA molecules, which means that the di-electrophoretic effect becomes apparent because the DNA molecules can be polarised. Therefore the total velocity field now consist of the electroosmotic, electrophoretic, and dielectrophoretic velocity fields, which can be written as:

$$\vec{u}_{DNA} = \mu_{EOF/EP}\vec{E} + \mu_{DEP}\vec{\nabla}E^2 \tag{2.84}$$

where $\mu_{EOF/EP}$ is the combined electroosmotic and electrophoretic mobility and μ_{DEP} is the dielectrophoretic mobility [43]. The device that is used for these experiments is the same shape, however the height of the device is now 400 nanometres. In total, two types of DNA molecules are used; λ-DNA with a contour length of 48 kbp (kilo base pairs) and T4GT7-DNA with a contour length of 165.6 kbp. The purpose of the experiment is to investigate what influence the DNA size and the type of applied electric field has on the pathlines of the DNA molecules. This will help in the development of a continuous, sieve-less, size-based DNA separation device [143]. Please note that the (unconfined) radius-of-gyration of both the DNA molecule types (λ-DNA = 0.74 micrometre, T4GT7-DNA = 1.37 micrometre) are larger than the height of the device, therefore the molecules will be confined and squeezed by the upper and lower walls of the device. The strength of the dielectrophoresis, and thus the ability to change the pathlines of the DNA molecules, depends on the gradient of the square of the electric field. In the analysis of EOF and EP flow it was established the strongest electric field (gradients) are located where the straight channel inlet is connected to the semi-circular chamber, therefore this area will be the centre of attention in the following

Table 2.5: List of parameter choices in the experiments, $V = a\sin(2\pi ft) + b$, along with the respective dataset names and colours used in the plot

Parameters / DNA molecule	λ-DNA (48.5 kbp)	T4GT7-DNA (165.6 kbp)
$a = 0$ V, $f = 0$ Hz, $b = 7.5$ V	Dataset A (black)	Dataset F (grey)
$a = 0$ V, $f = 0$ Hz, $b = 15.3$ V	Dataset B (black)	Dataset G (grey)
$a = 14$ V, $f = 1$ Hz, $b = 15.3$ V	Dataset C (black)	Dataset H (grey)
$a = 14$ V, $f = 1$ kHz, $b = 15.3$ V	Dataset D (black)	Dataset I (grey)
$a = 14$ V, $f = 1$ MHz, $b = 15.3$ V	Dataset E (black)	Dataset J (grey)

experiments and simulation results of the electric field intensity. However, the numerical mesh first needs some more attention. Because infinitely sharp corners cannot be practically achieved during the fabrication, and the fact that *if* sharp corners would be simulated it would result in a non-realistic singularity of the electric field intensity at that corner, the channel corners in the simulation are (slightly) rounded. However, it must be noted that this does not significantly change the qualitative distribution of the electric field intensity in the whole device, but only results in a more smooth distribution near those corners.

The bottom part of figure 2.20 shows the numerical results of the the electric field intensity in the area of interest, obtained using the finite element method. The top part of figure 2.20 shows two examples of the experimentally obtained pathlines of DNA molecules driven by an electric field of $V = 15.3$ V (DC), where the top left figure shows the pathlines of λ-DNA molecules, while the top right figure shows the pathlines of T4GT7-DNA molecules. It is very difficult to compare and analyse the two experimental images using only visual inspection. For example, qualitative comparison between the results obtained without DEP show very similar pathlines, even though quantitatively this may not be the case. Therefore some image analysis steps on the measured fluorescent images are performed using the image processing toolbox DIPimage [122]. For this purpose first a "background image" by averaging all images (typical number of images is ≈ 1000) in one particular measurement is created. The "background image" is then subtracted from each of the images in that measurement in order to remove unwanted signals such as auto-fluorescence from the channel walls and static objects in the channels. Afterwards, DNA molecules in each resulting images are segmented from the background using a fixed threshold algorithm [215]. The analysis is only done in a region-of-interest (ROI), which is defined as a rectangle around a single DNA molecule of interest, where the position of the molecule is determined by measuring the center-of-gravity across the ROI. In order to track the moving molecule, the ROI is also moved along with the molecule. This is done by using the measured molecule's center-of-gravity to specify the center of the moved ROI. Individual pathlines are then determined by tracking the molecules through the image sequence in the measured data.

The next step in the analysis is to quantitatively compare and analyse each pathline in all measurements and for that purpose two parameters as shown in the left figure of figure 2.21 are defined. First the "start distance", r_{start}, as the distance along the x-axis from the first corner in the walls is measured. After this, the "finish distance", r_{finish}, as the distance along the y-axis from the second corner in the same walls is measured. The distances (from pixels) are then calculated using the known scale of the photo (to micrometres). The electric field applied during the measurements can be expressed by the equation $V = a\sin(2\pi ft)+b$, where f and a are the frequency and amplitude of the AC signal, respectively, while b is the DC offset superimposed on the AC signal. The total contour length of the channels between the inlet and outlet ports, where the electric fields are applied, is approximately 2 cm. Table 2.5 shows the parameters chosen for the experiments, resulting in 10 datasets. In order to test the repeatability of the observations, a second set of experiments more than one month after the first set of experiments is performed, using different microchips. The plots in figure 2.21 and figure 2.22 summarise the results of the experiments, where the notations "1" and "2" on the datasets are used to indicate the first and second set of experiments, respectively.

In Figure 2.21, r_{finish} vs r_{start} for all datasets are plotted. If size-based DNA sorting occurs, then λ-DNA and T4GT7-DNA will systematically have different pathlines, and consequently have different values of r_{finish}, even though they start at the same position in the inlet and have the same value of r_{start}. Next the black markers (datasets A-E, i.e. λ-DNA in 400-nm

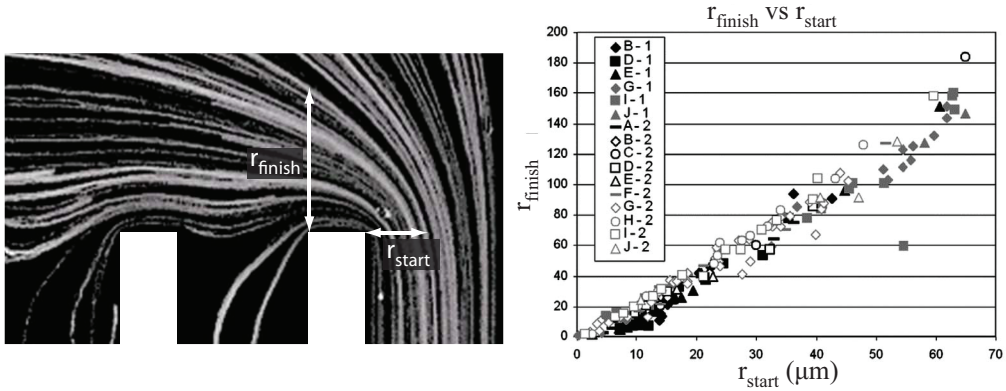

Figure 2.21: An illustration of the two parameters, r_{start} and r_{finish}, used to analyse the experimental data and the plot of r_{finish} vs r_{start} for all datasets.

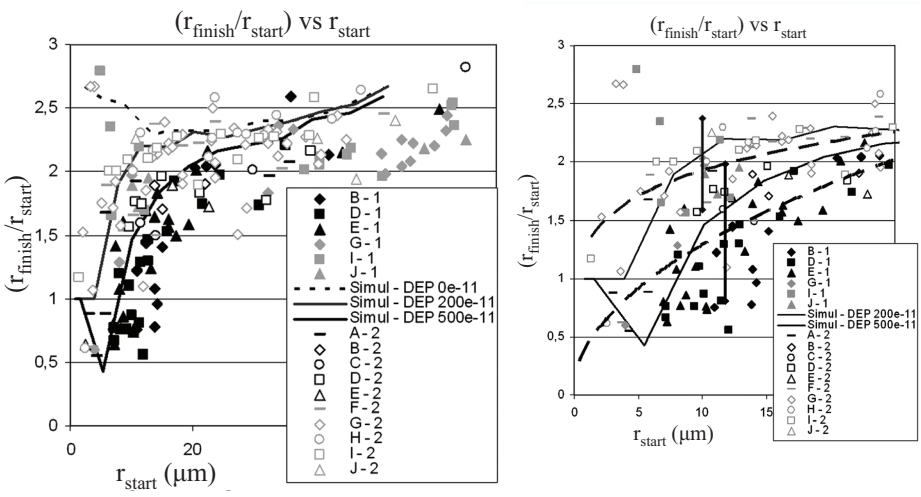

Figure 2.22: The plot of (r_{finish}/r_{start}) vs r_{start} along with the simulated pathlines, for all r_{start} values in the left figure and a close-up for $r_{start} < 25$ in the right figure. The dashed lines and error-bars represent trend lines and error margins of the experimental data (only for visual aid).

channels) and the grey markers (datasets F-J, i.e. T4GT7-DNA in 400-nm channels) in figure 2.21 are compared. Particularly for $r_{start} <\sim 25$ μm, the black markers consistently have r_{finish} values lower than the grey markers. Physically this indicates that a mix of λ-DNA and T4GT7-DNA in 400-nm deep channels can be sorted for $r_{start} <\sim 25$ μm, because for the same values of r_{start} their pathlines end up with different r_{finish}. Meanwhile, figure 2.22 show the plots of (r_{finish}/r_{start}) vs r_{start} for the same datasets. Compared to the plot in Figure 2.21, these plots allow a better assessment of the trends of the datasets. Again, for $r_{start} < \tilde{}25$ μm the black markers have lower values of (r_{finish}/r_{start}) than the grey markers. This confirms the remark above that for $r_{start} <\sim 25$ μm a mix of λ-DNA and T4GT7-DNA inside the channels can be sorted.

To compare the experimental results with the numerical results, the electrokinetic DNA pathlines are also computed using the finite element method. In the simulation, the ratio of dielectrophoretic mobility, μ_{DEP}, and the combined electroosmotic and electrophoretic mobility, $\mu_{EOF/EP}$, as defined in equation 2.84 are changed to match the experimental

results. This ratio is noted as, $\rho_\mu = \mu_{DEP}/\mu_{EOF/EP}$ and has the units $^{m \cdot kg}/_{s^7 A^3}$. The simulation results are plotted in figure 2.22, for $\rho_\mu = 0$ (i.e. when there is no dielectrophoresis effect in the system; shown as black box with lines), $\rho_\mu = 200 \cdot 10^{-11}$ (shown as grey line), and $\rho_\mu = 500 \cdot 10^{-11}$(shown as black line); the latter two simulation cases represent positive dielectrophoresis, where the DNA electrokinetic pathlines are shifted toward the highest electric field gradient indicated in the numerical results shown in figure 2.20. Comparison between the simulation and experimental data in figure 2.22 shows that the simulation of $\rho_\mu = 200 \cdot 10^{-11}$ matches the datasets with black markers, while the simulation of $\rho_\mu = 500 \cdot 10^{-11}$ matches the datasets with grey markers (In the right figure in figure 2.22, also the trend lines and error margins of the experimental data are shown for visual aid). This shows that the pathlines of the larger DNA (T4GT7-DNA) are attracted by dielectrophoresis toward the walls corner less strongly than the pathlines of the smaller DNA (λ-DNA). These observations are surprising, because Chou et al. reported that, in contrast to the above observations, dielectrophoretic effects on DNA molecules increase when the DNA size is increased [38]. A possible explanation is that as the DNA molecules pass the sharp corners at the U-turn, almost all of the monomers in the shorter DNA (λ-DNA) are affected by the attracting dielectrophoretic force while only some monomers in the longer DNA (T4GT7-DNA) are affected by the same force. The dielectrophoresis effects are therefore rendered to be less effective on the longer DNA than on the shorter DNA. On top of that, we also have to note that the experiments above are done in a nanofluidic setup, while Chou et al. used a microfluidic setup which has a much smaller confinement effect on the DNA molecules.

In order to see the effect of modifying the applied electric fields, the trends of datasets in each colour of the markers (i.e. comparison between datasets A-E among the black markers, and between datasets F-J among the grey markers) are compared. The plots in figures 2.21 and 2.22 show that there is no clear distinction between the trends of datasets in each colour of the markers. In other words, no significant effect by changing the applied electric fields in the experiment can be observed. However, the AC fields in the experiments only have a frequency of 1 Hz, 1 kHz or 1 MHz. Further investigations could also be done with higher frequencies, because the dielectrophoresis effect has been reported to increase when the AC frequency is increased [38]. Meanwhile, the results from the experiments using DC fields also indicate that DNA molecules experience positive dielectrophoresis, even though there is no periodic modulation of the electric field as typically caused by AC fields. This could be explained by the highly non-uniform ion and counter-ion concentrations along the nanofluidic channel depth, particularly when the depth is sufficiently small such that the electric double layer from the upper and lower channel walls are very close to each other, which is the case here. This would cause significant dynamic transverse transport of the ions along the channel depth (i.e. the z-axis) [147], and consequently the combination of DNA molecules and their surrounding counter-ions are perpetually polarized along the channel depth (i.e. the z-axis) [15], even though no AC fields are applied between the channel inlet and outlet. Any significant electric field gradient, such as caused by the sharp corner at the U-turn in the nano-device, would therefore cause positive dielectrophoresis to the DNA molecules regardless of whether the applied electric fields are AC or DC.

This section showed the numerical computations and experimental results of electroosmotic, electrophoretic and di-electrophoretic flows obtained inside nano-sized devices. Especially, it was shown that DNA molecules can be sorted based on their size particularly for $r_{start} <\sim 25$ μm, i.e. for pathlines that are close enough to the highest gradient of the electric field intensity in the device. This indicates that by the combination of electroosmotic, electrophoretic, and di-electrophoretic flow, for the first time, size-based DNA sorting can be done in a continuous, sieve-less, manner. The sorting is enabled by dielectrophoretic forces, which alter the DNA pathlines. These results are novel compared to the previous results of other authors, where dielectrophoresis has been used either for continuous sorting but only for cells or for spherical beads [43], or for manipulating DNA but only by trapping and not by continuous sorting [38, 157]. Meanwhile, for $r_{start} >\sim 25$ μm, no sorting can effectively be done because the pathlines are too far away from the corners in the channel walls, hence no significant dielectrophoresis effect exists. Please note however, that the dimensions and configurations of the channels in the branched U-turn nanofluidic channels are not yet optimised for an effective sorting operation and several geometry modification could be done to the sorting of DNA more effective [143]. However, it is also shown that the numerical simula-

tion, which where based on the pure continuum equations for electrokinetic flow, neglecting the fact the underlying phenomenon is basically on the molecular scale, can successfully be used to simulate the nanofluidics in this device. Furthermore, the numerical simulations also indicate that the observed pathlines can be explained by dielectrophoresis, which occurs even under DC electrical fields, presumably due to the nanofluidic confinement.

2.4.2 MD: Investigating Flow Characteristics Inside Nano-sized Channels

In the previous section it was shown how several nano-scale phenomena that are basically the result of molecular behaviour can be described satisfactorily by pure continuum equations. In this section the opposite is done: here (pure) molecular dynamics simulations on nano-sized channels are performed, and any deviations from continuum mechanics are explored.

With MD, the geometry used in the simulation must be constructed using atoms. Therefore, the nano-sized channel is simulated using two parallel solid atomic walls placed at a certain distance from each other, while between the two walls, atoms representing a liquid are placed. Figure 2.23 demonstrates this setup in 2D and 3D projections. The total number of atoms simulated is $N = 2048$, while the simulated (average) particle number density is, $\rho = 0.8$ units. Each wall of the nano channel consists of two layers of solid atoms, placed in a fcc lattice, while the total number of atoms in the liquid is $N_l = 1536$. The centers of the two walls are separated by 11.9698 units, which corresponds to approximately 4.1 nanometres when using the parameters for argon (see Table 2.3). The two walls will act as a boundary condition for the MD domain in the z-direction. Periodic boundary conditions are specified for the remaining boundaries of the MD domain. The MD domain is kept at a constant temperature, $T = 1.2$, using the Nosé-Hoover thermostat [72]. During each simulation several properties are binned. The MD domain is divided into 500 bins in z-direction and each bin collects the local value of the density, velocity, and the six stresses: Σ_{xx}, Σ_{yy}, Σ_{zz}, Σ_{xy}, Σ_{xz}, and Σ_{yz}.

In the first simulation the values of density, velocity, and stresses in the bins are collected, while *no flow* (i.e. no applied bodyforce) is present. The MD simulation is run for a total of 6.4×10^6 timesteps, which is just over 69 nanoseconds when using LJ-parameters to convert to physical units. The averaged values of each bin are computed by averaging over the last 5×10^6 time steps of the simulation. As expected, the collected (average) velocity is zero with some small fluctuations caused by the discrete behaviour of MD. However, the other results need closer attention. Figure 2.24 shows the results for the density, shear stress Σ_{xz}, and the averaged normal components of the stress, $P = \frac{1}{3}(\Sigma_{xx} + \Sigma_{yy} + \Sigma_{zz})$, referred to as the pressure inside the nanochannel. The density profile shows large variations around the average value of $\rho = 0.8$ units, where the variations near the wall are the largest. These variations are the result of the interaction of the atoms in the liquid with the solid atoms of the wall and can also be observed experimentally [34]. The variations are not a result of the flow, because no flow is present in this case. Further, Bitsanis et al. [21] showed that indeed the flow itself does have negligible effect on the density profile even when high shear rates are simulated. However, the large variations are present because of the finite temperature of the MD system, because this means that atoms are always in motion and therefore interacting with the wall and the other atoms all the time. The result is that atoms in the liquid spend more time inside certain layers near the wall, where the interaction between the wall and the surrounding liquid is more in equilibrium than elsewhere. These layers are the peaks visible in the density profile. On the other hand, because of the dense layer of atoms, the remaining atoms can not come close to this layer, because of the strong repulsion at close range. That means that on either side of a dense layer fewer atoms are present on average, which are the valleys in the density profile. Only when the atoms in the liquid are far enough from the wall, the interaction between the surrounding atoms is much more important, and the result is that, on average, or more or less constant density is reached. The results for this very small nano channel do not show a constant density anywhere and Other researchers [84, 177] showed how this is more evident in even smaller channels. Therefore, any flow behaviour of this nano channel or smaller nano channels, could not be computed with a standard continuum method. The same applies to the pressure profile, which is shown in the right part of Figure 2.24. This profile has large variations near the wall, while the values in the middle show fewer variations. Even without flow, the profile still shows a finite value,

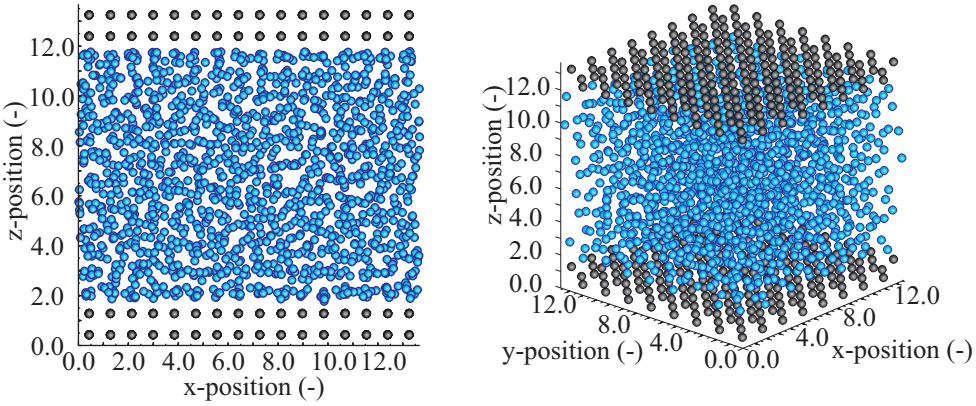

Figure 2.23: The geometry used for the MD simulation of flow in a nanochannel. The cubic MD system box is ~ 12 units which is ~ 4.1 nanometres. The total number of simulated fluid atoms are $N_l = 1536$. The MD system has periodic boundary conditions in x and y-direction, while the z-direction has an atomistic wall as boundary condition.

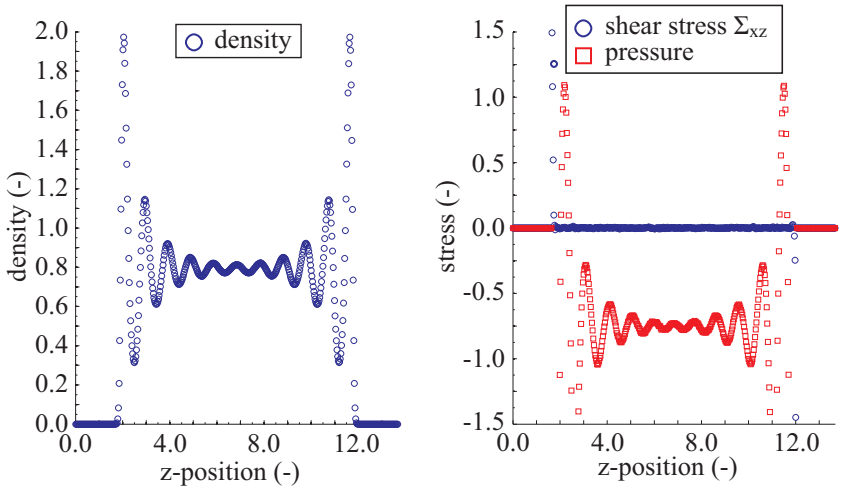

Figure 2.24: The density profile (left) and stress profiles (right) inside the nanochannel when there is no flow present. The profiles are obtained by binning 5×10^6 time steps.

because it is also determined by the interparticle forces between the atoms. On the other hand, the shear stress is zero, as is expected in the absence of a net flow.

In the next simulation, the effect of the wall interaction on the flow field is investigated. When using a continuum method, flow is generally applied by either (1) applying a pressure difference between the two sides of the channel, (2) by specifying the velocity itself at the boundary, or (3) by applying a (volume) body force. The first two options in the case of MD are not straightforward, and even so, the periodic boundary conditions will limit the applicability anyway. Therefore, the option that remains, is to apply a body force, which in MD means that each atom undergoes an additional force. Of course, when applying a body force inside a continuum fluid where the no-slip condition applies at the walls, the flow field will be that of a Poiseuille flow and the velocity profile is parabolic. However, as shown in the previous paragraph and shown by other researchers using Poiseuille flow [106, 107, 194, 195], deviations from this result can be expected in the case of a MD nanochannel. This MD simulation is run for a total of 2.4×10^6 timesteps, and the average values of the bins are collected from the last 1.6×10^6 time steps. Figure 2.25 shows the resulting density, velocity, shear stress, and pressure profile along the z-direction in the nano channel (illustrated in figure 2.23) when a bodyforce of $f_b = 0.100$ units is applied. The density profile in Figure 2.25 is almost exactly the same as the density profile in the absence of flow, Figure 2.24. This result accentuates the fact that the variations of density are a result of the wall-fluid molecular interaction, rather than the flow of the liquid. The overall velocity profile resembles a Poiseuille flow, however the velocity does show slight variations, especially near the wall. This is of course a direct consequence of the variations in density. When the density is higher, more atoms are present at a certain location, and because of this their motion is impaired. So, again the result show clear non-continuum aspects. This effect is especially visible in even smaller channels. For example, Travis et al. [194, 195] showed that for systems where the height is less than 4 units (less than 1.3 nanometres in the case of argon) the velocity profile is no longer parabolic. However, in the next section it will be discussed how much (or little) the velocity profile deviates from that of a continuum Poiseuille flow in the nano channel simulated here.

The following two results are the pressure and the shear stress. The pressure in the fluid with a bodyforce in Figure 2.25 is almost the same as the pressure in Figure 2.24. The only difference is the pressure due to the actual flow field, which is negligible to the total pressure that arises from the interparticle forces. The shear stress on the other hand shows a completely different result. In the continuum case, the shear stress profile in a Poiseuille flow is equal to a straight line, where the absolute values at the wall are the highest, and the shear stress is exactly zero in the middle of the channel. It is very obvious that the same qualitative behaviour is found with the MD simulation, where only near the wall, variations are clearly noticeable. In the next section this fact is used to determine a value for the viscosity of the liquid.

The previous simulation showed that when a certain bodyforce is applied to each atom of the liquid, the velocity profile closely resembles that of a velocity profile expected for a continuum with a constant viscosity, namely Poiseuille flow. However, the walls in the previous simulation were ideal, i.e. exactly lined in a fcc lattice and without surface defects. The next simulation simulates a nanochannel where the walls are not exactly in a fcc lattice and where the surface of the walls is not perfect smooth anymore. This basically simulates (real life) roughness of a nanochannel and the purpose of the simulation is to find the resulting velocity profile. Figure 2.26 shows the geometry that is used for this simulation. The MD domain is three times longer than previously used, making the total number of atoms that is simulated, $N = 6144$. The (periodic) roughness is modelled by adding several imperfections on the surface. Also every atom inside the wall is moved away from the perfect fcc lattice by adding a certain sinusoidal function to each of the atom's position. This results in the wavy character of the wall shown in the figure. The same bodyforce as before is applied, i.e. $f_b = 0.100$, and the other simulation parameters remain unchanged. The distribution and number of bins is adapted. Now two sets are defined: one set for the flow visualisation and one set for general post-processing. The first set collects the average values in 300 $\times$ 100 bins in the x-direction and z-direction respectively. The second set of bins collects the average values in 8 $\times$ 500 bins, which gives detailed profiles in the z-direction at 8 different position in the x-direction inside the channel. The simulation is run for a total of 2.4×10^6 timesteps, and the averaged values of the bins are collected from the

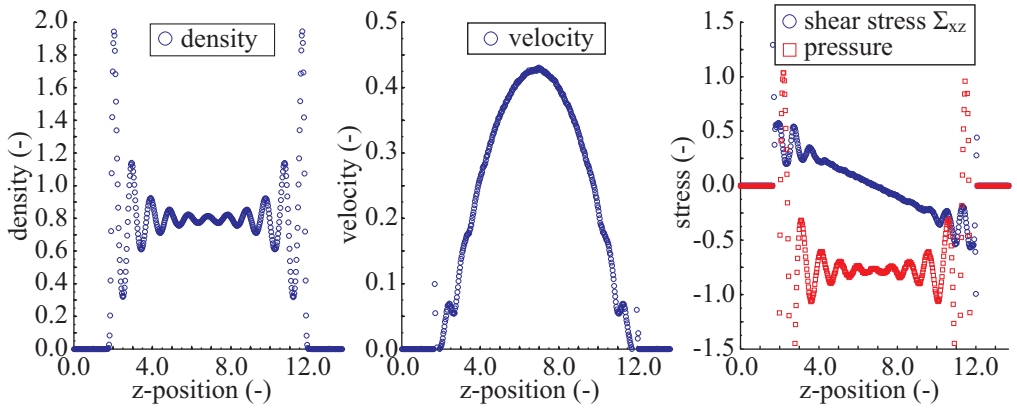

Figure 2.25: The density profile (left), velocity profile (middle), and stress profiles (right) inside the nanochannel when the bodyforce, $f_b = 0.100$. The results are obtained by binning the last 1.6×10^6 time steps (~ 17.2 nanoseconds). The velocity profile is very much similar to a (Poiseuille flow) parabola, however near the wall deviations occur.

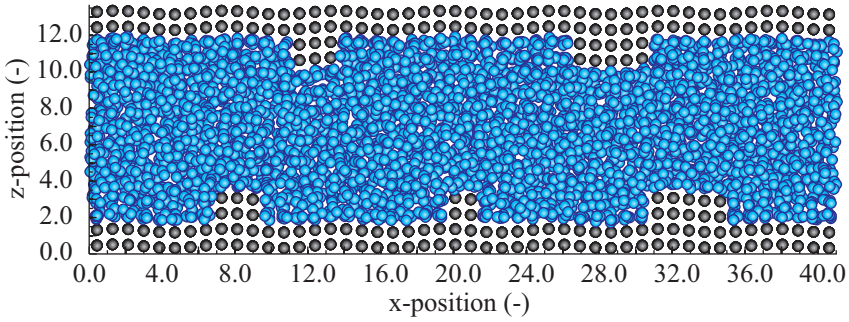

Figure 2.26: The bumpy geometry used for the nanochannel to simulate molecular rhoughness. This MD system is three times longer than the smooth channel while all other parameters are kept the same.

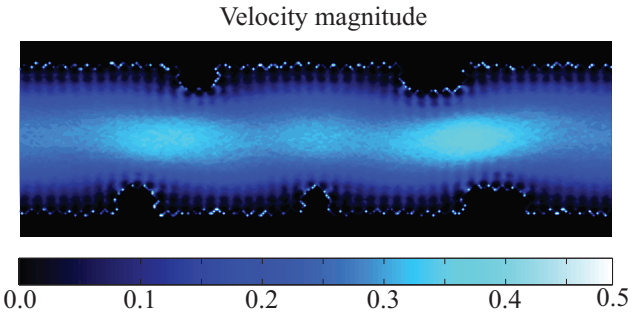

Figure 2.27: The velocity magnitude contours in the nanochannel with rhoughness.

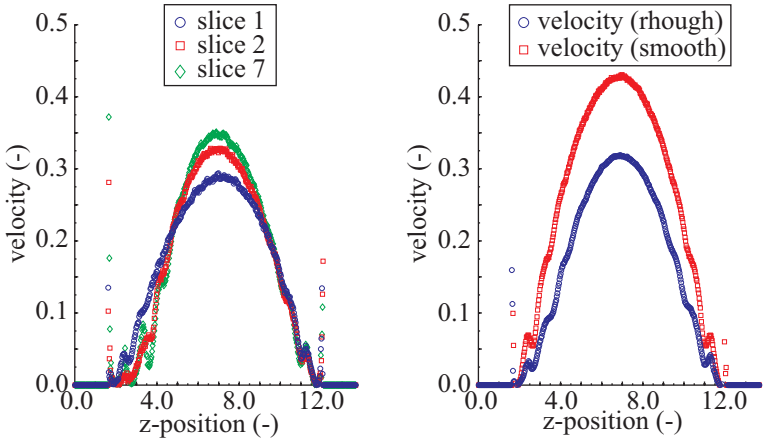

Figure 2.28: The velocity profile at three different positions in the nanochannel. The right figure shows a comparison of the average velocity profile of the nanochannel with rhoughness and the nanochannel with smooth walls, with the same bodyforce applied, $f_b = 0.100$ units. In both cases the velocity profile is similar to a parabola, except for near the wall. However, equivalent to the continuum case, the maximum velocity is less in the channel with rhoughness than that with the channel with smooth walls.

last 1.6×10^6 time steps.

Figure 2.27 shows the resulting velocity magnitude contours. Three peaks in the flow are visible, coinciding with the most narrow cross-sections due to the imperfections at the wall. The flow pattern basically corresponds to what can be expected for a continuum channel with constrictions, however the peaks in the figure are now less than a nanometre in size. Near and around the imperfections on the wall the variations in velocity, in other words the layering of atoms, is again visible. The results of the MD simulation show a very detailed view on the velocity in the nano channel, which is generally not available in experiments. An experimentalist observing the flow in this nano channel will at best see the average velocity profile. This is where the second set of the bins is used for.

Figure 2.28 shows several velocity profiles that are collected from the second set of bins. The left figure shows the velocity profiles in the z-direction taken from bins 1, 2, and 7, in the x-direction, respectively. The right figure shows the average velocity profile along all 8 bins in the x-direction. This is the velocity profile that an experimentalist would likely observe. As a reference the velocity profile in the nanochannel without the wall imperfections (i.e. the smooth wall) is also shown. Both velocity profiles are the result of the flow inside the nanochannel with the same value of the body force. In other words, for an external observer, equivalent to the continuum case, it looks like the channel has roughness, and therefore the resulting maximum velocity is less than that with a smooth wall.

In the previous simulation, imperfections where added to the wall to simulate (actual) roughness. However, there are other ways of simulating wall-fluid behaviour. In the Lennard-Jones model two empirical parameters are used for the interparticle behaviour of the fluid atoms, i.e. ϵ and σ. However, it is possible to specify different values of ϵ and σ for the interaction between the wall and the fluid. In the next simulation the parameter ϵ_{wf}, which is the interaction strength between the wall and the fluid atoms, is changed from 1.0 to 0.7. As before, the simulation is run for a total of 2.4×10^6 time steps and the average values of the bins are collected from the last 1.6×10^6 time steps. The MD domain is divided into 500 bins in z-direction, and the flow is induced by a body force, $f_b = 0.100$ units.

Figure 2.29 shows the resulting velocity and shear stress profiles. As a reference, the velocity and shear stress profiles for the smooth nano channel with $\epsilon_{wf} = 1.0$ is shown. The velocity profile in the nano channel with the reduced wall-fluid interaction shows a shift of the velocity; the shape is in most parts the same, however an apparent wall slip occurs. This fact is more clear in the shear stress profile, where indeed the shear stress is exactly the same

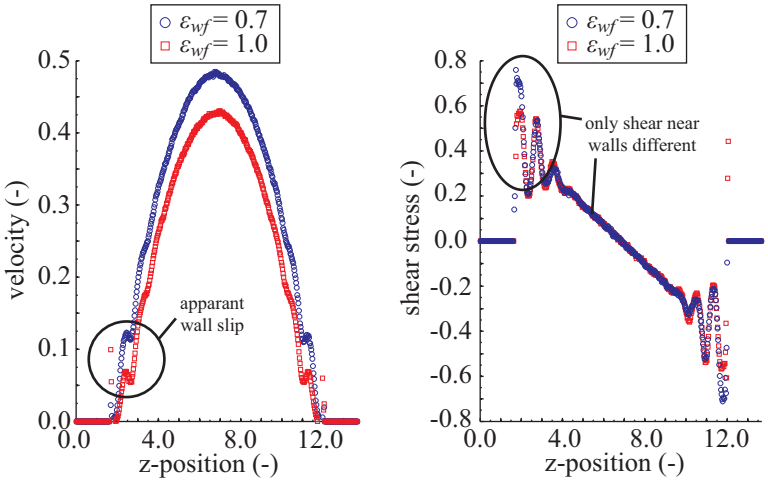

Figure 2.29: This figure shows how the velocity profile changes, if the parameter ϵ_{wf} is changed. Now a slip velocity near the wall is visible.

everywhere except very near the wall. There the absolute value of the shear stress is larger than the value for the nano channel with the strong interaction. The equivalent property in the case of a pure continuum computation would be the specification of a certain slip velocity. When a slip velocity is specified at the wall, the Poiseuille velocity profile simply shifts. In the case of the MD simulation, the reduced interaction between the wall and the fluid means that the fluid atoms are less influenced by the wall. Together with the body force this means that the fluid atoms, on average, are less slowed down when they interact with the wall. This last fact also means that the density profile changes slightly.

All results of the previous simulations show clearly that large variations in properties are present near the wall, and gradually become less at larger distances from the wall. Despite that, the previously simulated nano channels were too small to observe any averaged values without significant variations due to the finite number of atoms, i.e. they still showed continuum-like behaviour. However, in the next simulation a nano channel where the walls are further apart is simulated. This is done to investigate the minimum size of the channel needed to obtain "true" continuum properties. Just like before, each wall of the nano channel consists of two layers of solid atoms, placed in a fcc lattice. The centers of the two walls are now separated by 25.6496 units, which corresponds to approximately 8.2 nanometres when using the parameters for argon in Table 2.2. The total number of atoms simulated is $N = 4096$. Because the channel is higher, a smaller value for the body force is needed to achieve the same flow velocity. The used body force is, $f_b = 0.010$ units. The other simulation parameters remain identical.

Figure 2.30 shows the resulting density, velocity, pressure, and shear stress profiles in the larger nanochannel. Although there are still large variations visible near the wall, the values in the middle of the nano channel have almost no visible variation. The density profile shows that in the middle of the channel the averaged particle number density is 0.8, as expected. The averaged velocity profile shows a closer match to a (continuum) Poiseuille velocity profile than obtained by the previous simulations, while the same applies to the shear stress.

The results of the simulations show that in a nanochannel with a height of about 25.6 units, or 8 nanometres, yield results near that expected for a continuum. Also results obtained by, for example Travis et al. [195], suggest that the velocity profiles in nano channels with a height of 10 units show very small deviations between continuum and MD results. However, this does not yet mean that a pure continuum method to compute the flow inside this channel would be advisable. The variations in the values of several properties are noticeable until about 8.5 units distance away from *each* wall. For the nanochannel this means that a total of 17 units out of 25.6 units experiences variations in values of properties, which in effect means that

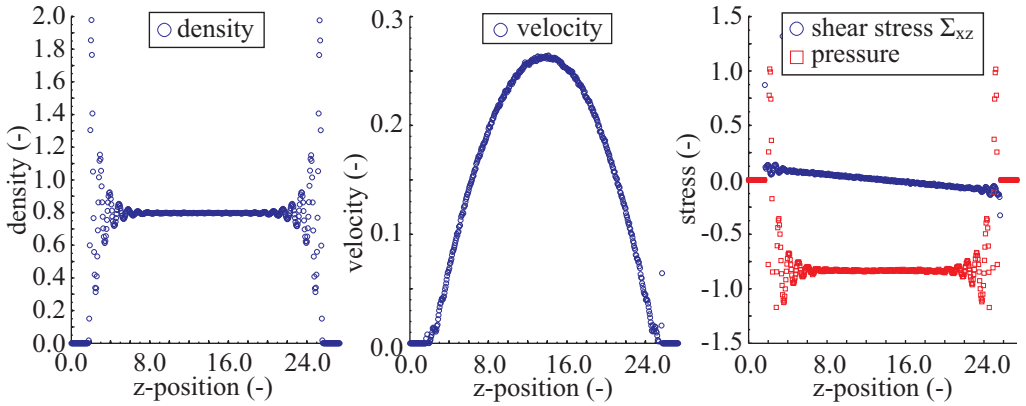

Figure 2.30: The density profile (left), the velocity profile (middle), and stress profiles (right) for a nanochannel where the walls are further apart. This channel has a height of ~ 25.6 units which is ~ 8.2 nanometres. The total number of atoms simulated is $N = 4096$, while the applied bodyforce, $f_b = 0.010$.

about 65% of the nanochannel experiences non-continuum effects. This information however gives a possibility to find the minimum required height of the nano channel before a pure continuum method provides reliable results. For example, the continuum computation of a nano channel with a height of 170 units, which is approximately 58 nanometres, neglects only 10% non-continuum effects.

This section showed how flow in a nano channel can be simulated with MD, and how changing several parameters can affect this flow. However, all simulations had in common that near the atomistic wall large variations in values can be observed, while away from the wall these variations become smaller. Although these non-continuum effects are noticeable, the velocity profile still resembles that of a Poiseuille flow. In the larger nanochannel the variations in the middle of the channel even disappear, and basically motion of the atoms can be considered as a continuum. Especially for larger channels the MD method is really only required near the wall, while the rest of the channel can be computed with a much more efficient method than MD. One way of doing this is discussed in Chapter 3.

2.4.3 MD: Obtaining Viscosity Using Poiseuille and Couette Flow

In the previous section several results from simulations of flow inside a nanochannel were presented. However, MD is more than just a flow simulation technique; by simulating a certain amount of atoms in a specific ensemble, important material properties can be obtained. In this section one of such material properties, namely viscosity, is obtained from such a simulation.

There are several ways how viscosity can be obtained from a MD simulation. The Green–Kubo relations [76, 112] are commonly used, which is an exact expression for linear transport coefficients in terms of integrals of time correlation functions. The value of the viscosity in these cases are predominately taken from a simulation that simulates the bulk, e.g. using an infinite periodic system. However, there are other possibilities, which are generally faster. In the previous section it was shown how a Poiseuille-like flow was generated using an externally applied force, f_b, inside a nano channel. This gives the possibility to relate the resulting Poiseuille flow to the viscosity of the liquid, which is argon in this case. However, Bitsanis et al. [22] showed that the effective viscosity inside the channel can increase considerably in very narrow channels, generally smaller than 4 units. However, when the size of the channels is increased, the bulk value can be obtained. This is in line what was noted before; that, except for near the wall, continuum-like behaviour can be observed whenever the channel is large enough. Therefore the bulk viscosity can be obtained from the velocity profile inside this channel using the equation relating the viscosity and the velocity inside

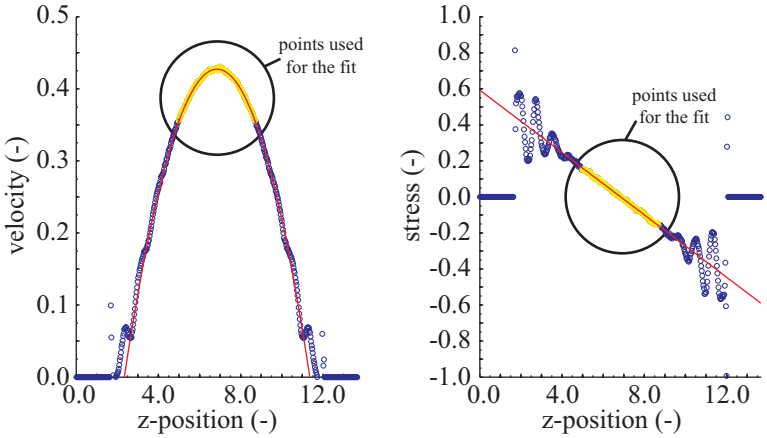

Figure 2.31: The viscosity can be determined from a Poiseuille flow using two different methods. The first method fits the velocity profile and compares it to a pure Poiseuille velocity profile (left), while the second method uses the shear stress and velocity gradient to obtain the value of the viscosity (right). The selection of the values that are used for the fit are based on where the shear stress profile does not show any apparent variations (see text for details).

the channel, which is given by:

$$\frac{\partial}{\partial z}\left(\mu\frac{\partial u}{\partial z}\right) = -\rho f_b \tag{2.85}$$

In order to determine the viscosity from this equation, the second derivative of the velocity profile needs to be measured. The easiest way to do this is to curve fit the resulting velocity profile from the MD simulation. From this curve fit the second derivative can be taken. An alternative relation that can be used to determine the viscosity is using the relation between the shear stress and the velocity gradient perpendicular to the flow:

$$\Sigma_{xz} = \mu\frac{\partial u}{\partial z} \tag{2.86}$$

In this case, the first derivative from the curve fit of the velocity profile is taken and the value for the shear stress is binned during the simulation and is therefore known. However, the evaluation of the curve fit needs further attention. The curve fit is trying to fit a Poiseuille velocity profile (a second order polynomial fit), or in other words, a continuum profile. Near the wall there are large variations that are non-continuum. Therefore the curve fit must include only those values that are continuum-like. Figure 2.31 shows the same velocity profile and shear stress profile displayed in Figure 2.25, but now includes the collection of values that are used for the fit (different colour) and the fit itself (solid line). The selection of the values that can be used for the fit are based on the shear stress profile. Only those points on the curve are used that correspond to a straight line, i.e. no variation are apparent for these values. The same datapoints are used for the velocity profile fit. The fitted curves which corresponds to a continuum profile, can now be used to determine the viscosity using 2.85, i.e. method 1 or 2.86, i.e. method 2. Table 2.6 shows the viscosity determined from various simulations, with different applied body force. The amount of time steps over which the values were binned is shown in the second column. When the body force is small, the simulation runs over a longer time in order to collect more statistics to distinguish the average value from the discrete noise. The third and fourth columns show the value for the viscosity obtained using 2.85 and 2.86, respectively. The errors in these values are computed using the 95% confidence bounds obtained from the fitted curves. Although the values for the viscosity fall nicely within the range of values obtained by others [106], there is some variation visible in the obtained values for the viscosity and the errors. In general, the errors of the fit obtained with method 2 are larger. The predominant reason for this is the fact that more statistics need to be collected to obtain meaningful values for stress inside MD simulation than for velocity, i.e. the value of stress is similar to the noise of

Table 2.6: Obtained viscosity from the simulation of the nanochannel where the total number of atoms, $N = 2048$ and $T = 1.2$. The results are obtained by using two different methods (see text).

f_b	timesteps binned	μ method 1 (-)	μ method 2 (-)
0.005	18.4×10^6	1.840±0.07	2.054±0.42
0.010	18.4×10^6	2.035±0.04	2.166±0.22
0.020	8.8×10^6	1.950±0.03	2.099±0.17
0.050	3.2×10^6	1.988±0.02	2.147±0.11
0.100	1.6×10^6	2.013±0.02	2.091±0.09
0.200	1.6×10^6	2.210±0.01	2.143±0.04

Table 2.7: Obtained viscosity from the simulation of a nanochannel where the walls are further apart. The total number of atoms, $N = 4096$.

f_b	timesteps binned	μ method 1 (-)	μ method 2 (-)
0.002	6.4×10^6	2.053±0.02	2.209±0.15
0.005	6.4×10^6	1.983±0.01	2.138±0.08
0.010	3.2×10^6	1.975±0.01	2.133±0.06
0.020	1.6×10^6	1.972±0.01	2.126±0.04

the MD simulation. Another observation is the fact that the value of the viscosity obtained with method 1 seems to rise with increasing body force. The reason for this is most likely caused by the increasing shear rates inside the very small channel. However, the very small size of the nano channel simulated here could also be a factor. This is discussed next.

The nanochannel simulated previously is a very narrow channel, about 4.1 nanometres in height. As described before, the non-continuum effects are present inside the entire nanochannel. These effects may affect the value of the measured viscosity. Therefore another set of simulations is performed with a wider nanochannel. The height of the nano channel is doubled to about 8.2 nanometres, where the total number of atoms simulated becomes $N = 4096$. The results of these simulations are shown in Table 2.7. The values for the viscosity are similar to those of the previous simulation, but more consistent than before. Therefore it can be concluded that in the middle of the nanochannel indeed continuum-like behaviour occurs. However, note that the (local) viscosity near the atomistic wall changes dramatically and is not equal to the value of viscosity in the bulk of the nanochannel.

So far only Poiseuille flow was simulated. However, another type of flow that is easy to deal with, is Couette flow. Couette flow occurs as the flow between two no-slip walls, where one or both walls are movin g. The main advantage is that no body force needs to be applied. It is the no-slip boundary condition that implies that the velocity of the fluid is the same as the velocity of the wall. This principle can also be applied to MD. The atoms in the atomistic wall are moved with a certain constant velocity. When the wall atoms move they will start dragging the fluid atoms with them. In order to determine the viscosity of the MD fluid two simple simulations are performed. Again the nanochannel with $N = 2048$ is used, where at the top the wall atoms are given a velocity of 0.1 and 0.2, respectively, while at the bottom the wall atoms are give the velocity of -0.1 and -0.2, respectively. In each simulation the shear stress is also binned.

Figure 2.32 shows the results from the simulation. As expected, the shear stress is nearly a constant in the middle of the channel, while near the walls large variations are again seen. From these results, the viscosity can be determined using 2.86. The viscosity determined using the curve "Couette 1" is $\mu = 2.084 \pm 0.11$, while the viscosity determined using the curve "Couette 2" is $\mu = 2.098 \pm 0.05$. These values are again similar to the results obtained previously and reported in Tables 2.6-2.7.

In this section it was shown how molecular dynamics can be used to determine the viscosity of the simulated liquid inside a nanochannel. In this case argon was simulated and the viscosity is approximately 2.0. This result is very important for simulations where MD and

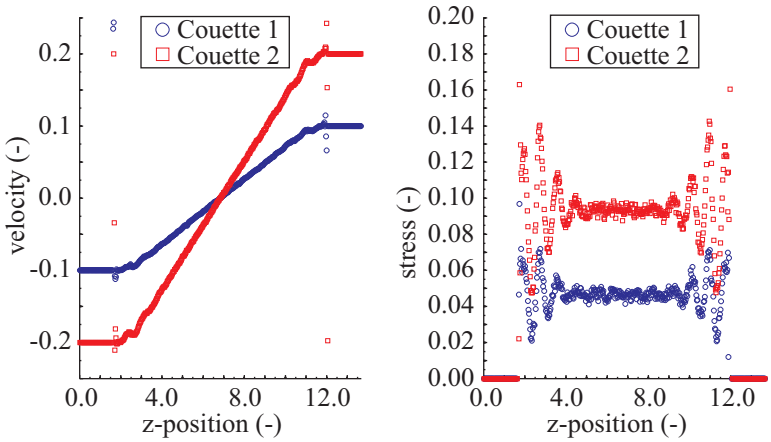

Figure 2.32: The viscosity of argon is determined for a Couette flow. This method uses the shear stress and velocity gradient to obtain the value of the viscosity. The total number of atoms, $N = 2048$ and the temperature, $T = 1.2$.

the continuum become coupled. Unlike what is the case for MD, where the viscosity is a result of the simulation, the continuum flow equation requires a value for the viscosity in order to solve the equation.

2.4.4 MD: Using Particle Mesh Methods to Accurately Simulate Surface Tension

The simulations so far used a cutoff radius of $r_c = 2.5$ to optimise the computation. In Section 2.3.2 it was argued that in this case the interaction energy is only 0.8% of the maximum attraction value possible. It is therefore reasonably safe to neglect all molecular interaction beyond the cutoff radius and to correct the final result of the simulation (e.g., a desired material property) with a long-tail correction (if needed). However, there are exceptions, and one important one is encountered in the study of multiphase flows, where the effect of surface tension is crucial. For example, Nijmeijer et al. [138] showed how the value of the surface tension in the simulation is increased by a factor of 2.8 if the cutoff radius is increased to $r_c = 7.33\sigma$ and Holcomb et al. [90] and Meche et al. [130] concluded that a cutoff radius of $r_c = 5.0\sigma$ together with a tail correction is required to obtain a reliable value for the surface tension. Trokhymchuk et al. [196] confirmed that the value of surface tension still changes up to 10% when the cutoff radius is increased from $r_c = 4.4\sigma$ to $r_c = 5.5\sigma$. Also, more recently Sinha et al. [175] found that truncating long-range terms of the Lennard-Jones potential function at $r_c = 4.5\sigma$ would cause errors as high as 15 percent in surface tension of argon, but also showed how this can be eliminated by using the PPPM method for long range terms for the $1/r^6$ term inside the Lennard-Jones potential. Bo Shi et al. [172] showed how the same method can also be used to accurately simulate the surface tension of water. Very recently [171], the same authors also showed how the contact angle of a Lennard-Jones liquid and water droplets adjacent to a solid surface can be predicted in good agreement with experimental values, using the same long-range technique.

However, another interesting phenomenon is the nano-jet and nano-jet breakup, which has application in future inkjet printing. Most of the models and theories used for the simulation of inkjet printing do use continuum models, however because of the small scale of the nanojet, these models give different results than a pure MD simulation predicts. On the other hand, a pure MD simulation can also give the possibility to accurately and easily model wall-fluid interactions like wetting, hydrophilic or hydrophobic behaviour. Recently a lot of interest is shown in the simulation of the break-up of the (long) nanojet [36, 45, 55, 133, 173, 190], where because of the small scale not only the surface tension, viscous, and intertial forces play a role in breakup, but also the thermal fluctuations in the interfacial region. These fluctuations also tend to suppress the formation of satellite drops, i.e. smaller drops than

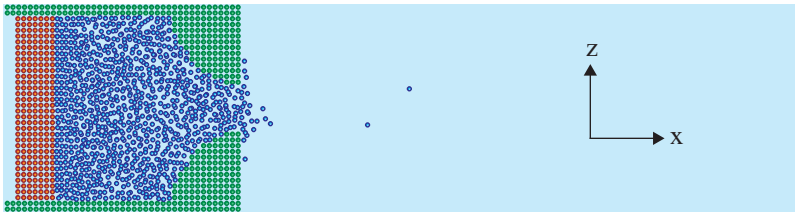

Figure 2.33: The MD model of the nano inkjet printer used for the simulations. The liquid is contained inside a reservoir which is connected to a nozzle. The liquid is set into motion by a moving wall at the left. This snapshot is taken just after the moving wall (red atoms) is set into motion.

the main drop, which is a desirable feature in inkjet printing. Molecular dynamics can also be used to study important practical parameters of the inkjet printing device, e.g. the final droplet size. Choi et al. [36] showed what the influence of the nozzle outlet size is on the breakup time and on the growth rate of spherical droplets. Dai et al. [45] investigated the whole nanojet ejection process for various compressing velocities. Their results showed large variations in density and pressure during the ejection process and a linear relationship between the length of the liquid threads and the compressing velocity was found. However, connected to what was mentioned in the previous paragraph, it is interesting to investigate the influence of the used cutoff radius on the actual droplet formation process and the final droplet. Shin et al. [173] used a cutoff radius of $r_c = 3.5\sigma$, while Choi et al. [36] used $r_c = 3.0\sigma$ and Dai et al. [45] $r_c = 2.5\sigma$. Clearly these values are lower than the value of the cutoff radius where correct surface tension values are obtained. In this section a simulation of a nano-sized inkjet printer and the break-up of a nano-jet is investigated. More specifically, the method where the PPPM method is used to determine the $1/r^6$ term inside the Lennard-Jones potential is used to compare the creation of the droplet and the break-up of the nanojet, obtained with different cutoff radii.

The atomistic model that is used for the simulation is depicted in Figure 2.33. The device consists of a 9.6×6.6 nanometre rectangular reservoir where the left side of the reservoir is mobile and can push out the liquid. The right side of the reservoir is connected to a nozzle that concentrates the liquid (argon) towards the nozzle outlet with a 3.5 nanometre width. The walls of the device are created using atoms that are placed in a FCC lattice where the total number of wall atoms is 10912. The number density of the liquid is $\rho_N = 0.8$, which means that the total number of atoms corresponding to the liquid inside the device is 13920 atoms. The boundary conditions of the MD domain are periodic in the y-direction, while the boundary condition in the x-direction and z-direction is a specular wall that keeps the atoms inside the MD domain. The total MD domain size is approximately $45.3 \times 9.3 \times 10.5$ nanometres. The fluid atoms interact with each other according to a Lennard-Jones potential, although the wall-fluid interaction needs further attention. The hydrophobic or hydrophylic capabilities of the surface can influence the ability to create droplets and for this purpose the interaction potential of the wall and fluid atoms can be modified [136]:

$$U_{wall-fluid}(r_{ij}) = 4\alpha_{wf}\epsilon\left[\left(\frac{\sigma}{r_{ij}}\right)^{12} - \beta_{wf}\left(\frac{\sigma}{r_{ij}}\right)^{6}\right] \tag{2.87}$$

where the parameters α_{wf} and β_{wf} are the hydrophylic and hydrophobic interaction factors, respectively. High values of α_{wf} and β_{wf} ($\alpha > 1$ and $\beta > 1$) will result in a strong interaction, which means that it simulates a hydrophylic surface, while low values will result in reduced interfacial resistance, simulating a hydrophobic surface. Before the droplet creation process is initiated by moving the left side of the reservoir chamber, the whole MD domain is equilibrated and thermostatted with a Nosé-Hoover thermostat to a temperature of $T = 0.9$. During the equilibration the liquid atoms are kept inside the device by means of a specular wall at the nozzle outlet to prevent atoms escaping, i.e. evaporate. This specular wall is removed once the equilibration is achieved. At that point, the wall at the left side is set into motion with a certain velocity. This wall keeps on moving at this velocity over a certain distance, and then it is stopped. This is done for two reasons: The first reason is that

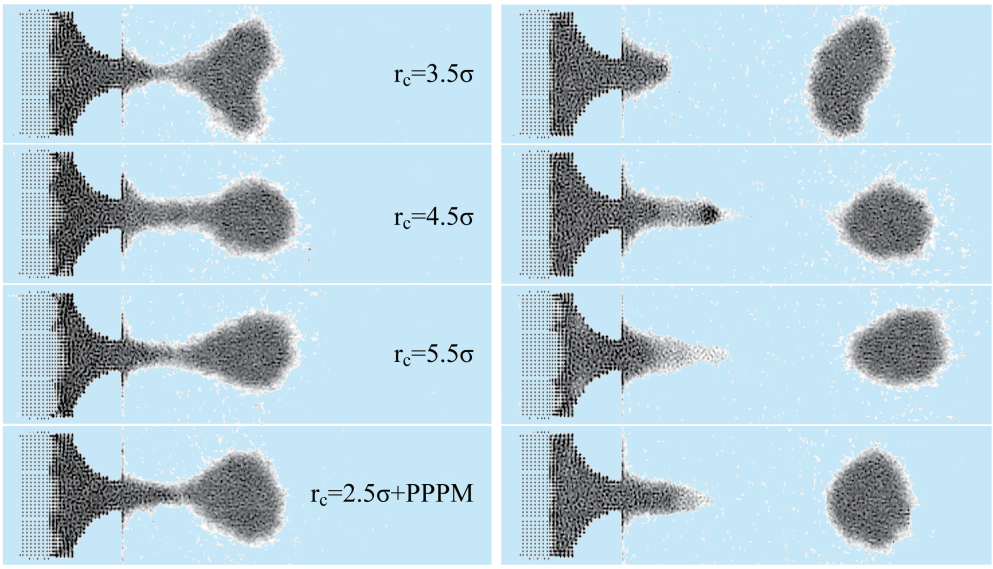

Figure 2.34: The simulation of the creation of a droplet is sensitive to the employed cutoff radius. The left sequence (top to botttom) shows how the computation of the tail enlongation change when the cutoff radii $r_c = 3.5\sigma$, $r_c = 4.5\sigma$, $r_c = 5.5\sigma$, and the long-range method for the van der Waals interaction with a cutoff radius, $r_c = 2.5\sigma$ are used. The same is shown for the droplet creation in the right figure. Note that the results for every sequence are shown at the same point in time of the simulation.

during the motion a certain amount of inertia builts up, and when the wall stops moving, this energy is large enough to overcome the energy associated with the surface tension to create a droplet with a certain size and velocity. The second reason is that the displaced volume of fluid inside the chamber is nearly identical to the volume of the droplet that is created. This gives a possibility to control the droplet size.

The first set of simulations is performed to investigate the influence of the cutoff radius on the numerical results of the droplet creation. The cutoff radii that are investigated are $r_c = 3.5\sigma$, $r_c = 4.5\sigma$, and $r_c = 5.5\sigma$. The results from these simulations is also compared to a simulation where the van der Waals attraction force is obtained using the PPPM long-range technique, while the cutoff for the other interactions is $r_c = 2.5\sigma$. However, the MD system only has periodic boundary conditions specified in *one* dimension, while the other dimensions are specified by a specular wall. This means that the PPPM-method must be adapted to prevent wrong summation in the non-periodic directions. In section 2.3.6 several methods where specified that deal with this kind of systems, however the method used here is the 2D-Slab method with an added (vacuum) space at both sides of the inkjet device and includes the ELC term to prevent the slab-slab interaction in z-direction. This means that the long-range interaction in x-direction is (wrongly) considered periodic, however it is expected that the simulation results are not severely changed by this. The grid size for the PPPM method is $[256 \times 32 \times 96]$ with splitting parameter $\beta = 0.790$. The velocity of the moving wall is $v_{wall} = 0.50$ and it is moved inwards over a distance of about 4.5 nanometres and then stopped. In this simulation, the parameters used for the wall-fluid interaction are, $\alpha_{wf} = 1.0$ and $\beta_{wf} = 1.0$, which means that the interaction is identical to the fluid-fluid interaction. Figure 2.34 shows the comparison of the formation of the droplet (left) and the resulting drop (right) simulated with the different cutoff radii, where the snapshots are all taken at the same point in time. The results obtained from the simulation where the cutoff radius, $r_c = 3.5\sigma$, shows an unrealistic droplet formation and resulting drop because the surface tension of the liquid is underestimated which results in non-spherical shapes. On the other hand, the results obtained from the simulation where the cutoff radius, $r_c = 4.5\sigma$ and $r_c = 5.5\sigma$, show much more realistic shapes, while the difference between $r_c = 4.5\sigma$ and $r_c = 5.5\sigma$ are only visible in the details of the droplet break-up. Finally, the simulation

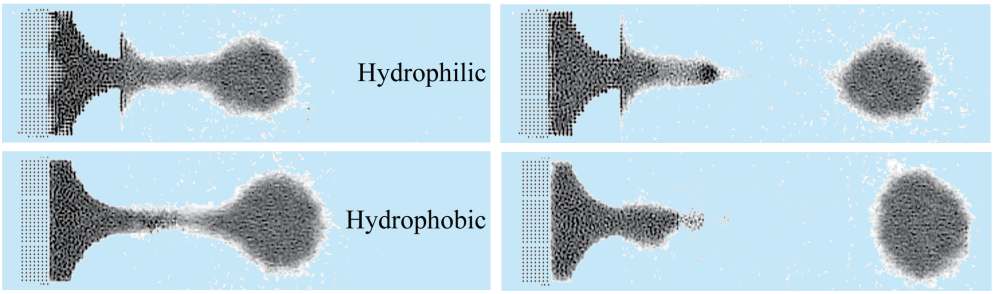

Figure 2.35: This figure demonstrates how the simulation of the creation of a droplet is sensitive to the hydrophilic or hydrophobic characteristics of the wall.

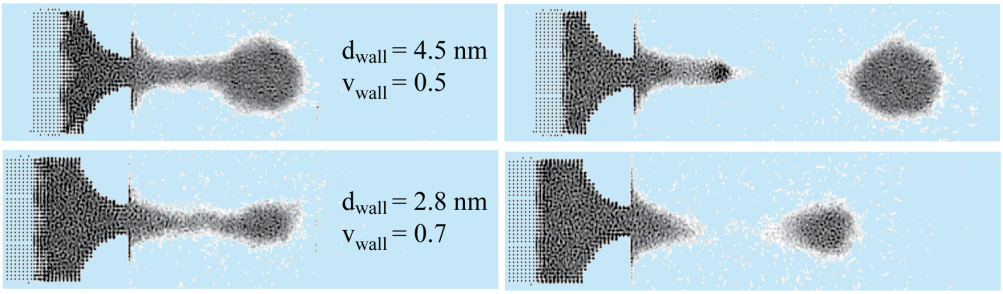

Figure 2.36: This figure demonstrates how the simulation of the creation of a droplet is sensitive to the total distance the moving wall is moved and how the velocity of the moving wall needs to be changed in order to create a droplet.

with $r_c = 2.5\sigma$, but where the $1/r^6$ term inside the potential is obtained using the long-range PPPM technique, show results which are very close to those obtained where all the interaction is obtained using large cutoff radii. Besides droplet size and shape, the velocity of the droplet is also an important parameter for inkjet applications. The velocity of the droplet is easily obtained from the simulation results and in all cases the value was between $v_{drop} = 0.5$ and $v_{drop} = 0.6$, indicating that this parameter is not highly sensitive to the applied cutoff radius.

The second simulation is performed to investigate the effect of the hydrophobic or hydrophylic properties of the surface. For this simulation the parameters inside the wall-fluid interaction potential are taken as $\alpha_{wf} = 0.3$ and $\beta_{wf} = 0.4$, which means that the surface now is strongly hydrophobic. The employed cutoff radius is $r_c = 4.5\sigma$, while the other simulation parameters are the same as before. Figure 2.35 shows the comparison between the simulation with $r_c = 4.5\sigma$ with the hydrophylic surface done before and the new simulation with the hydrophobic surface. As expected, the big difference is the fact that in case of a hydrophobic surface, the liquid does not want to be in contact with the outside of the nozzle and therefore more liquid is pushed into the tail of the droplet and the droplet itself, which also results into a bigger droplet size in the end. It also has effect on the velocity of the droplet, which now is $v_{drop} = 0.8$ and is higher than when a hydrophylic surface was simulated.

The third simulation is performed to investigate what effect the velocity of the moving wall and the total distance that the wall moves has on the resulting droplet formation and droplet size. For this simulation, the wall is only moved to the right over a distance of about 2.8 nanometres, which is 1.7 nanometres less than before, while the velocity of the wall is kept the same at $v_{wall} = 0.50$. However, the result of this simulation showed that the built-up energy was not large enough to overcome the energy involved in creating a droplet and therefore no droplet was created and all the liquid collected near the nozzle entrance due to surface tension. Therefore another simulation was performed where the velocity of the wall

was increased to $v_{wall} = 0.70$. Figure 2.36 shows the results of this simulation, where now a much smaller droplet is created. The velocity of this droplet is however also much smaller, $v_{drop} = 0.35$.

In this section a nano-sized inkjet device was investigated, especially the influence of the cutoff radius on the resulting droplet formation. It was shown that the long-range technique where the van der Waals interaction is computed using a PPPM technique while the rest of the interaction uses a cutoff radius of $r_c = 2.5\sigma$, shows similar results compared to simulation without the long-range technique but with much larger cutoff radii. Also, the employed cutoff radius does have a big impact on the actual computed droplet creation process, e.g. the elongation of jet and ultimate shape and size of the droplet. Surprisingly, the cutoff radius did not seem to have a lot of effect on the computed droplet velocity, which is an important model parameter for inkjet printers. On the other hand, it was shown that the wall-fluid interaction can have a major impact on both the droplet size and velocity. However, it was also demonstrated how the droplet size can be changed by changing the distance the wall of the inkjet device is moved.

2.4.5 MD: Comparing the Value of Viscosity for Several Water Models[†]

Viscosity plays an important role in many physical processes and hence it is important to specify it accurately in computer simulations in which these processes are investigated. Of special interest is water. When simulating water with molecular dynamics (MD) it is vital to choose a water model that correctly predicts the process investigated. Many water models have been developed, differing in parameter values and number of charge sites, and each having different succes in predicting the correct value of a certain physical parameter. However, only few references give a complete set of the viscosity versus temperature of a certain water model, which makes it difficult to choose the appropriate water model. In this study, the viscosity temperature relation of four water models is considered. These are the popular 3-point charge SPC/E water model [16], and several variants of the 4-point charge models, TIP4P [102], TIP4P/Ew [92], and the recent TIP4P/2005 [2].

There are several ways to obtain the values of the viscosity by means of MD simulations and most of them require many statistics to be collected. The most frequently used methods are the Green-Kubo method [76, 112] and the Stokes-Einstein method [57]. Both methods are based on the autocorrelation function of the pressure tensor, which is computationally expensive to obtain. Another way of finding the viscosity is by simulating Couette shear flow, which for argon was done in section 2.4.3, however, this still requires the computation of the ratio of pressure tensor and strain rate. The quality of viscosity obtained from these methods strongly depends on the accuracy of the pressure, which in turn strongly depends on the used cut-off radius and method used to compute the long-range interactions [87]. Alternatively, the periodic perturbation method [75] can be used, where the viscosity can be calculated from a steady-state velocity profile generated by a periodic external force applied to the system. Compared to the pressure, the velocity profile is straightforward to extract from a MD simulation with less statistics that needs to be collected to obtain meaningful results.

Previous calculations of the viscosity for different water models showed that there can be a large difference between the calculated value and the experimental value, as discussed next. Most research concentrated on the SPC/E water model. For example Balasubramanian et al. [14] showed how the viscosity of SPC/E is calculated for $T = 303.15K$, using both equilibrium and non-equilibrium molecular dynamics. The value they found is about 18% less than the experimental value, which is similar to the error that Gou et al. [80] found. Hess et al. [87] used periodic shear flow to calculate the viscosity of the SPC/E water model at $T = 300K$ and found a value about 30% lower than the experimental value. Recently, this has also been verified by Chen et al. [35] using the Green-Kubo method. Wu et al. [211] simulated shear flow at $T = 298.5K$ and found an error of about 23% for SPC/E. On the other hand, Bordat et al. [23], using a reverse non-equilibrium molecular dynamics simulation, calculated the value of the viscosity for the SPC/E water model at $T = 300K$ which was almost the

[†]Submitted as: A P. Markesteijn, R.M. Hartkamp, S. Luding, and J. Westerweel, "A comparison of the value of viscosity for several water models using Poiseuille flow in a nano channel"

same as the experimental value (within 5% error). Less research is done on the other water models. For example, Yongli et al. [214] calculated the value of the viscosity at several liquid water temperatures for several water models (including the TIP4P model) using the Stokes-Einstein relation and reported errors between 30.3% and 52.3% between experimental values and calculated values. Wensink et al. [205] calculated the value of the viscosity for the TIP4P model at $T = 298.25K$ and found an error of more than 60% compared to the experimental value. Very recently, Song et al. [178] performed non-equilibrium molecular dynamics simulations using the periodic perturbation method to simulate the shear viscosity of five commonly used water models (including the SPC/E and TIP4P model). The value they found for the viscosity of the SPC/E model at $T = 300K$ is 15% less than the experimental value, while the value for the TIP4P model was 41% less than the experimental value. For the TIP4P/Ew water model no data for the viscosity is known to us, while for the TIP4P/2005 only very recently [73] the dependency of the value of the viscosity on the pressure at three different temperatures were calculated using the Green-Kubo method. The conclusion was that, at least at these temperatures, the value of the viscosity is very well predicted (slightly less than 5% error) with the TIP4P/2005 water model. Furthermore, in other recent papers [1, 7, 149] it was shown how the TIP4P/2005 water model predicts a wide range of material properties accurately and therefore is a promising water model.

In this section, several of the water models (SPC/E, TIP4P, TIP4P/Ew, and TIP4P/2005) will be tested on their ability to model the viscosity of liquid water between the temperatures $T = 273K - 373K$. Especially, it will be shown how Poiseuille flow generated inside a nano-sized channel can be used to extract the value of viscosity versus the temperature for these models very efficiently. The flow is generated by a constant body force on each of the water molecules inside the channel and the resulting velocity profile is used to calculate the viscosity. This technique was used before to investigate the viscosity of simple fluids, like argon, in Section 2.4.3, and has several benefits compared to the other methods. For example, the reported simulation times needed in order to obtain meaningful statistics to calculate the viscosity using the Green-Kubo method, the Stokes-Einstein method, or the Couette shear flow method are 10, 20, or 60 ns. Water simulations commonly use a time step of 1 or 2 fs, meaning that several tens of million time steps are required in these simulations and therefore require considerable computational effort. Although the periodic perturbation method performs better, where only 2 ns - 4 ns of simulation time is needed to obtain the results, here it will be shown that good statistics for the velocity profile can be obtained within only 1 ns of steady state flow. However, similar to the periodic perturbation method, it means that the viscosity is not obtained in a shear-free situation. Therefore, care must be taken that the shear inside the nano channel does not become too large.

The value of the viscosity is calculated for the different water models by examining Poiseuille flow inside a nano-sized channel illustrated in Figure 2.37. The Poiseuille flow is generated by a constant body force f_{bx} in the x-direction, to each molecule. The channel itself is created by modelling two parallel solid atomistic walls at a certain distance from each other in the z-direction. Between the two walls, the water molecules are placed. The boundary conditions in the x and y-direction are periodic, while in the z-direction the water molecules are constricted by the walls. The equation for ideal Poiseuille flow inside a channel in this case is:

$$\frac{d}{dz}\left(\mu\frac{du_x}{dz}\right) = -\rho f_{bx} \tag{2.88}$$

where u_x is the (macroscopic) velocity in the x-direction inside the channel, ρ is the liquid density, and μ is the unknown viscosity. The velocity profile and density can be extracted from the MD simulation, while the applied force is known. This gives the possibility to relate the resulting Poiseuille flow to the viscosity of the used water model. However, care must be taken, for example Bitsanis et al. [22] showed that, at least for simple liquids like argon, the effective viscosity inside the channel can increase considerably in very narrow channels (height smaller than 5 molecular diameters). One reason for this is the wall-fluid interaction which results in a layering effect of atoms near the wall and this effect only gradually disappears away from the wall (see Section 2.4.2). However, despite this, approximately quadratic velocity profiles can be obtained for simple liquids confined to channels only 10 molecular diameters in height [195]. For more complex liquids, like water, which is simulated in this study, the same phenomenon can be expected.

In order to extract the viscosity from the velocity profile, it is important that only the part

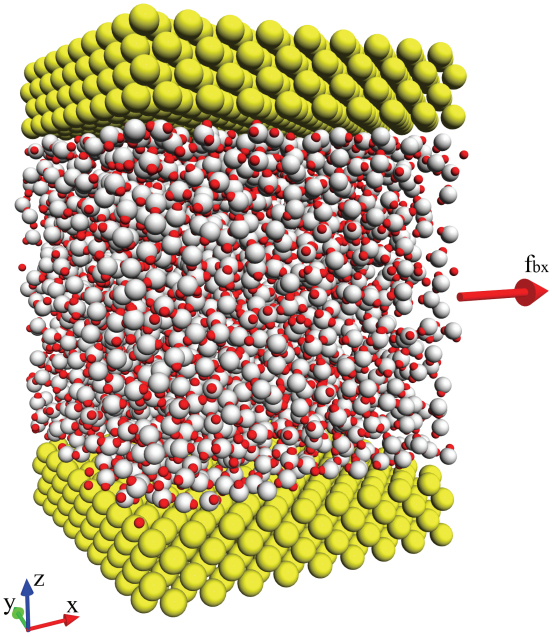

Figure 2.37: Figures showing the MD model of the nano channel (left). In total 2048 water molecules are placed between two solid atomistic walls each consisting of 648 silicon atoms. The distance between the centres of the two walls is ≈ 4.3 nm. Poiseuille flow is generated by a body force f_{bx} in the x-direction.

of the results are used that show the expected continuum behaviour. The results of the simulations with argon showed a layering effect of atoms near the wall due to the wall-fluid interaction and the used Lennard-Jones potential parameters. The same can be expected from simulations with water near an atomistic wall. To examine this, an equilibrium MD simulation (i.e. *without* flow) is carried out, where the variations near the wall and extent of the variations are studied, especially the density and charge distribution profiles. The density profile should show a region with a constant value, i.e. the expected continuum value, in the middle of the channel, while layering of molecules can be visible near the wall. The charge distribution profile is of interest, because this indirectly gives information about the orientation of the water molecules. However, note that in this study, the wall atoms are deliberately not charged. This means that the well-known phenomenon of the electric double layer (see Section 2.1.2) is not taken into account and the resulting velocity profile is only due to Poiseuille flow.

The water models that are investigated all have in common that they solve for the following potential energy equation between molecules i and j:

$$U = 4\epsilon \left[\left(\frac{\sigma_{OO}}{r_{OO}} \right)^{12} - \left(\frac{\sigma_{OO}}{r_{OO}} \right)^{6} \right] + \sum_{\alpha=1}^{N} \sum_{\beta=1}^{N} \frac{q_{i\alpha} q_{j\beta}}{4\epsilon_0 r_{i\alpha,j\beta}} \tag{2.89}$$

Lennard-Jones interaction between the water molecules is only considered between the O-atoms of each molecule i and j, where r_{OO} is the distance between the two atoms. The Coulombic interaction of the water molecules is computed using a total of N charge sites associated to each water molecule, where $q_{i\alpha}$ is the charge of the α^{th} charge site of molecule i and ϵ_0 is the electrical permittivity of vacuum. The bond lengths and angles are fixed using the SHAKE algorithm [165]. The water models differ in the values of the parameters they use. Table 2.4 gives an overview of the used parameters in the four models compared here.

For each water model, several MD simulations are performed in a (canonical) NVT ensemble. The temperature in the system is controlled by a Nosé-Hoover thermostat [91, 139] and the

simulated temperatures range from 273 to 373 K. The walls of the nano channel are placed approximately 4.3 nm apart, measured from the centre of the bottom wall to the centre of the top wall. Each of the atomistic walls consists of four layers of solid atoms, placed and fixed in a fcc lattice square, i.e. in ABAB layers. The chosen material properties of the wall are based on silicon, which has a density of $\rho_{wall} = 2329\,\mathrm{kg/m^3}$ at the initilisation temperature of the MD system, $T_{init} = 293$ K. Each wall consists of 648 wall atoms, which interact with (only) the oxygen atoms according to the Lennard-Jones 12-6 potential. The Lennard-Jones parameters for the wall in the case of Silicon are: $\sigma_{wall-O} = \sigma_{Si-O} = 3.24$Å and $\epsilon_{wall-O} = \epsilon_{Si-O} = 1.274\mathrm{kJ/mol}$ [103], however, as will be revealed later, slightly different parameters for the wall are used in the simulations where the viscosity is obtained. The total number of water molecules between the two walls is 2048. These are also initialised in a fcc lattice, with an initial density of $\rho_{init} = 998.2\,\mathrm{kg/m^3}$, and are allowed to melt during the equilibration process [63], which takes 0.3 ns.

The value of the body force f_{bx} is chosen such that the typical value of the maximum velocity inside the channel never exceeds $20\,\mathrm{m/s}$. This was done to prevent the shear inside the nano channel becoming too large. This maximum velocity is selected since it provides good statistics of the velocity profile compared to the natural vibrations of the water molecules, while it was verified that the value of the viscosity was not greatly affected by the amount of generated shear in the channel by performing a simulation with different values for f_{bx}.

The integration of Newton's equation of motion is performed with the Verlet algorithm, while the electrostatic interactions are treated by the 'Particle-Particle and Particle-Mesh' (PPPM) method (see Section 2.3.4.2). However, the MD domain only has periodic boundary conditions specified in two dimensions, the PPPM-method must be adapted to prevent erroneous summation in the direction perpendicular to the walls. The method adopted here is the 2D-Slab method with an added (vacuum) space at both sides of the wall, and includes the ELC term to prevent slab-slab interactions (see section 2.3.4.3 for details). The grid size for the PPPM method is $[36 \times 36 \times 60]$ with a splitting parameter $\beta = 0.305$, while the cutoff value for the Lennard-Jones interaction is $r_c = 1.0$ nm, which is $\approx 3\sigma_{OO}$. For each MD simulation, the total simulation time is 1.5 ns, with a MD time step of 1.0 fs. All results are determined from the last 1.2 ns of the simulation and are obtained by binning the different macroscopic values in 500 bins, which are equally distributed across the z-direction.

First the equilibrium MD simulation is carried out in order to study the variations near the wall. The water model SPC/E is employed for this simulation, for which the parameters can be found in table 2.4. During the simulation, the density and charge distribution values are collected in the bins.

Figure 2.38 (top) shows the density profile while the bottom figure shows the charge distribution across the channel. The density profile shows large variations around a more or less constant value in the middle of the channel, where the variations near the wall are the largest. These variations are the result of the interaction of the water molecules with the solid atoms of the wall and can also be observed experimentally [34]. Water molecules spend more time inside certain layers parallel to the walls, where the interaction between the wall and the surrounding liquid is more in equilibrium than elsewhere. These layers are the peaks visible in the density profile and come from the fact that, due to the strong LJ repulsion, have a minimal distance in z-direction, but can be disordered (inside the layer) in x-y direction. On the other hand, because of the dense layer of water molecules, the remaining molecules can not come too close to this disordered layer because of the strong repulsion at close range. That means that on either side of a dense layer, fewer molecules are present on average, which are the troughs in the density profile. Only when the water molecules are far enough from the wall, the interaction between the surrounding water molecules becomes isotropic due to increased disorder, and the result is that more or less constant isotropic density is reached.

The charge distribution profile shows that the orientation of the water molecules is also constricted near the walls. The first strong positive peak of the profile indicates that more H-atoms than O-atoms can be found near the wall, while the first trough after the peak shows a strong negative charge indicating a layer of mostly O-atoms coinciding with the density profile. The second peak and trough show something similar. On the other hand, in the middle of the channel the average charge is zero, indicating random orientation of the water molecules. Note that the total charge averaged across the entire profile equals zero, e.g. the water inside the channel is neutral. Although in the simulation the electrokinetic

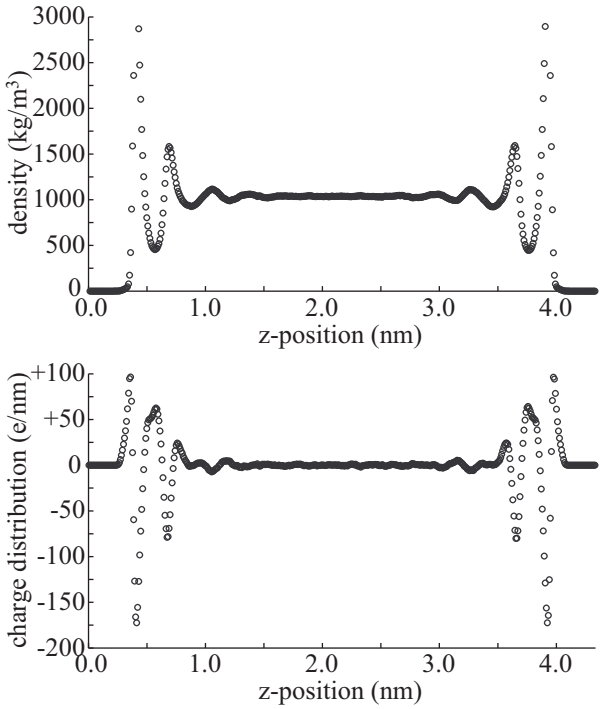

Figure 2.38: The density profile (top) and charge distribution profile (bottom) obtained for the nano channel from the equilibrium MD simulation ($\rho_{init} = 998.2\,\mathrm{kg/m^3}$, $T_{init} = 293$ K, using the SPC/E water model)

effect caused by charges of the wall and counter-ions in the water, i.e. the electric double layer, was deliberately not simulated, the orientation of the water molecule itself can create a charge distribution perpendicular to the wall.

Near bulk or continuum conditions are established 1.2 nm away from the walls, indicating the possibility to extract the value of the viscosity. However, the value of the density in the middle of the channel is about $1035\,\mathrm{kg/m^3}$, while the averaged density across the complete density profile is exactly the same as the initialised value, $\rho_{init} = 998.2\,\mathrm{kg/m^3}$. The reason for the higher value of density measured in the middle of the channel is the combination of the Lennard-Jones parameters for the wall-fluid interaction and the resulting interaction with the water molecules.

In order to compare the value of the viscosity of each water model to the theoretical value of the viscosity, each simulation should be performed with a predefined value of the density in the middle of the channel. The value of the density that is aimed for, is the value at 1 bar for the different temperatures, i.e. $\rho\,(T)|_{p=const}$ and can be found in for example Bird et al. [20]. There are several ways how this can be accomplished. One possibility is to change the distance between the two walls of the channel accordingly. This is comparable what is done in an NPT-simulation, where the pressure is kept constant by changing the size of the MD simulation box. However, considering the fact the simulations are done in an NVT ensemble, another way to simulate the correct density in the middle of the channel is to adjust the wall-fluid interaction parameter of the wall, σ_{wall-O}. If this value is decreased, water molecules are able to move closer to the wall and therefore, on average, spent more time in a larger volume between the two walls. This results in a lower density in the middle of the channel. The following simulations are therefore performed with a variable value of the wall-fluid interaction parameter σ_{wall-O}. The (macroscopic) density can easily be sampled and does not require a lot statistics to be collected before a meaningful value is obtained. Therefore, the correct value of σ_{wall-O} can be obtained within the equilibration process (i.e. within 0.3 ns of simulation time) and the remainder of the simulation is performed with this

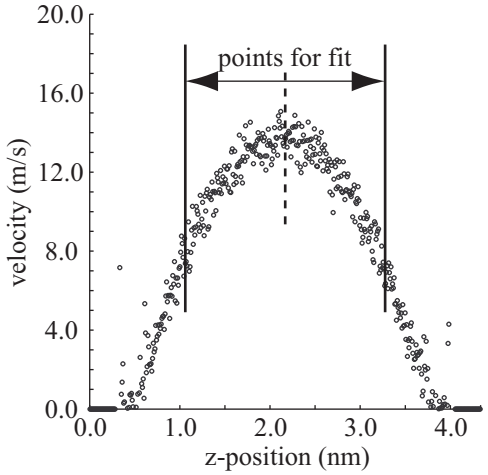

Figure 2.39: A typical velocity profile from one of the simulations of water inside a nanochannel (using TIP4P/2005). An illustration of the fraction of the velocity profile used for the curve fit is given.

Table 2.8: The value of viscosity of the four water models as obtained from the fitted velocity profile. The experimental values are given for reference.

T(K)	SPC/E μ(mPas)	TIP4P μ(mPas)	TIP4P/Ew μ(mPas)	TIP4P/2005 μ(mPas)	Experiment μ(mPas)
273	1.282 ± 0.0940	0.668 ± 0.0515	1.601 ± 0.1459	1.697 ± 0.1259	1.778
277	1.073 ± 0.0556	0.698 ± 0.0232	1.196 ± 0.0776	1.506 ± 0.1125	1.572
283	0.879 ± 0.0356	0.605 ± 0.0179	1.057 ± 0.0947	1.114 ± 0.0629	1.303
293	0.795 ± 0.0473	0.544 ± 0.0143	0.744 ± 0.0261	0.928 ± 0.0341	1.004
303	0.663 ± 0.0239	0.479 ± 0.0146	0.705 ± 0.0249	0.817 ± 0.0476	0.802
313	0.519 ± 0.0134	0.402 ± 0.0089	0.538 ± 0.0148	0.586 ± 0.0280	0.658
323	0.424 ± 0.0154	0.325 ± 0.0125	0.483 ± 0.0240	0.557 ± 0.0248	0.551
343	0.370 ± 0.0193	0.285 ± 0.0135	0.384 ± 0.0218	0.408 ± 0.0152	0.407
363	0.301 ± 0.0097	0.280 ± 0.0159	0.270 ± 0.0110	0.320 ± 0.0135	0.315
373	0.271 ± 0.0068	0.251 ± 0.0128	0.264 ± 0.0110	0.291 ± 0.0120	0.281

obtained value.

Next, the viscosity will be determined for the four different water models, SPC/E, TIP4P, TIP4P/Ew and TIP4P/2005 for a temperature range of $T = 273K - 373K$ by simulating Poiseuille flow in the nano channel. As noted before, the value of the wall-fluid interaction parameter of the wall, σ_{wall-O} is changed such that the density in the middle of the channel is equal to the (theoretical) value of the density of water at 1 bar. Furthermore, the interaction strength between the wall and the water molecules is changed/increased to: $\epsilon_{wall-O} = 3\epsilon_{Si-O} = 3.822^{\mathrm{kJ}}/_{\mathrm{mol}}$. This was done in order to reduce any significant slip developing near the wall.

Figure 2.39 shows a typical velocity profile obtained from such a simulation. As expected, the velocity profile is similar to a Poiseuille flow velocity profile, with only minor differences very near the wall. In order to determine the viscosity from equation 2.88, the second derivative of the velocity profile needs to be measured. The easiest way to do this is to curve fit the data points of the velocity profile from the MD simulation. From this *fitted* velocity profile, the second derivative with respect to the height of the channel can be taken easily. The data points are fitted to a function described by $u = u_0 + a\,(z - z_0)^2$, where the fitting parameters u_0 and a need to be determined and the parameter z_0 is taken as the middle of the channel. Doing so, means the obtained velocity profile is assumed to be symmetric,

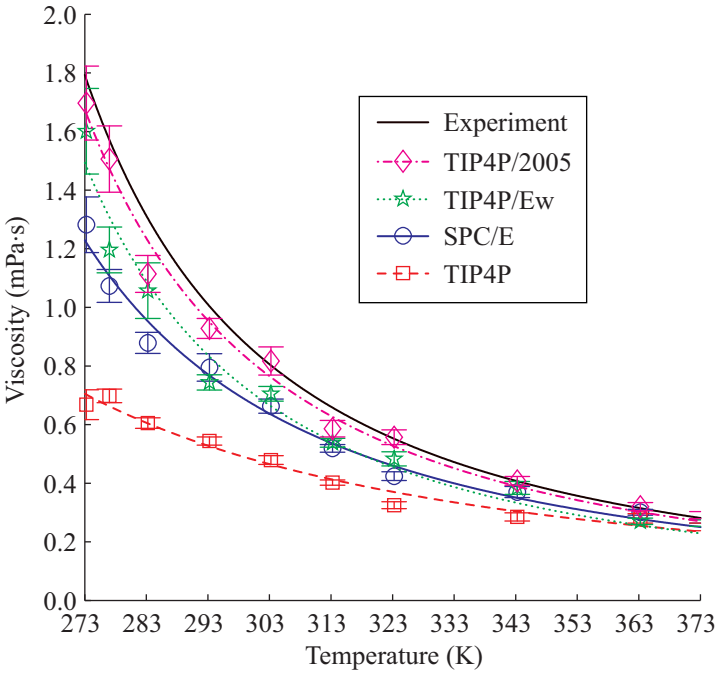

Figure 2.40: The values of the viscosity as a function of the temperature obtained from the curve fit of the velocity profile for the four different water models. The error of the fit is also displayed. The lines in the figure are obtained from a fit of the type: $\mu = (T - T_0)^{-b}$, where the fit with experimental data from Bird et al. [20] is used for reference.

Table 2.9: The parameters obtained from the fit of the type $\mu = (T - T_0)^{-b}$ from the datapoints as shown in figure 2.40.

	Experiment	TIP4P/2005	TIP4P/Ew	SPC/E	TIP4P
T_o	225.4	224.0	224.5	212.4	173.6
b	1.637	1.642	1.677	1.633	1.578

while the fitting parameter a is directly related to the second derivative and therefore the viscosity. However, because this value is sensitive to the actual fit, multiple symmetric curve fits of different (continuum-like) sections of the (same) velocity profile of the MD results are taken. The sections of the velocity profile that are used for the fit, are described by the fraction of the total velocity profile, where the value of 1.0 means the obtained velocity profile from wall to wall is used. In total 10 fits for each MD velocity profile are performed. The fractions for the set of curve fits are: 0.80, 0.75, 0.70, 0.65, 0.60, 0.55, 0.50, 0.45, 0.40, and 0.35. The final value for the viscosity is obtained by averaging the resulting set of fits and the error of the fitted value is estimated using the corresponding 95% confidence bounds of the set of fits.

Figure 2.40 shows the obtained values for the viscosity and the error of the fit for the four different water models. The experimental values of the viscosity are displayed for reference. The curves are obtained from a fit of the type: $\mu = (T - T_0)^{-b}$, fitted to experimental data from Bird et al. [20], where $T_0 = 225.4\,K$ and $b = 1.637$. Table 2.9 show the fitted parameters obtained for the simulation results of the four water models. Table 2.8 shows the obtained values of viscosity for the four different water models and the deviation from the theoretical value of viscosity in percent.

The results show that the TIP4P water model severely underpredicts the value of the vis-

cosity at all liquid water temperatures, especially at the lower temperatures, where the deviation from the theoretical value is 45% to 60%, as found also by others [205, 214]. The SPC/E model and the TIP4P/Ew water model show comparable performance in predicting the value of the viscosity. The TIP4P/Ew water model is slightly better, but does so with more computational effort because of the extra interaction site involved. In general the deviations are about 15% to 30% for the SPC/E water model and 10% to 25% for the TIP4P/Ew water model at the lower temperatures. The error for the SPC/E water model at $T = 293K$ and $T = 303K$ are very similar to the errors reported by other authors [14, 35, 80, 87, 211]. Both models predict the value for the viscosity within 10% to 15% for the higher temperatures. However, the TIP4P/2005 water model predicts the value of viscosity very accurately for the whole range of liquid water temperatures. The values are within about 5% to 10% from the experimental values of the viscosity.

Overall, the averaged data confirm the above conclusions; the TIP4P model underpredicts the viscosity by as much as -39%, the SPC/E and TIP4P/Ew water models underpredict the viscosity by 19% and 15%, respectively, while the TIP4P/2005 water model underpredicts the value of viscosity by (only) 3%.

This concludes this section where numerical results are presented that show how the viscosity for four different water models is calculated and compared to experimental values. This was accomplished by simulating Poiseuille flow in a MD nano channel. The value of the viscosity was determined by curve fitting the resulting velocity profile and comparison of the profile to a continuum solution. The benefit of using this method is the fact that good statistics for the velocity profile can be obtained within only 1 ns of steady state flow, which is considerably faster than alternative methods for finding the viscosity in MD simulations. The results from the simulations showed that the TIP4P/2005 water model is able to accurately predict the viscosity over a wide range of temperatures of liquid water, the TIP4P model performs the worst, while the SPC/E and TIP4P/Ew water model performed similar and resulted in reasonable accuracy for the value of the viscosity. Therefore, if a simulation is performed where the viscosity plays an important role, the TIP4P/2005 water model is recommended. Furthermore, in the next chapter the coupling of molecular dynamics and continuum is introduced. As will become evident, the value of the viscosity is a very important parameter for correct coupling of the two domains. Therefore, if a molecular dynamics simulation of water needs to be coupled with the continuum, the TIP4P/2005 water model is also recommended.

Chapter 3

Coupling MD and CFD: Argon and Water

In the previous chapter the two separate simulation techniques of computational fluid dynamics and molecular dynamics were discussed. Both numerical techniques use Newton's equations of motion to obtain the simulation variables like velocity and pressure. However, in MD the trajectories of individual atoms are followed that, when averaged, give the desired macroscopic variables, while CFD solves differential equations to obtain them directly. In this chapter a technique is explained that couples both numerical techniques, i.e. in one domain the CFD technique is applied while in another domain the MD technique is applied, these are coupled and the two region together will give an unique solution based on the two domains.

3.1 Schwarz Alternating Method

In Section 2.4 it was shown that using a MD simulation, basic continuum behaviour can be seen even at the very small scales. However, also large deviations to continuum behaviour can be observed, especially near the wall. In the region where there is continuum behaviour, a CFD technique could equally well be applied. The fact that the MD technique yields the same (continuum) solution, establishes a way to couple the two different regions.

When dealing only with partial differential equations (i.e. continuum description), it is well established that with domain decomposition methods a boundary-value problem can be solved by splitting it into smaller boundary value problems on sub domains. The solution to the whole domain can be obtained by alternatingly solving the same boundary value problem restricted to each individual subdomain. More specifically, in overlapping domain decomposition methods, the sub domains overlap by more than the boundary, which in practice means that the overlap is several grid cells. One of the earliest of these methods is the Schwarz alternating method [169]. This method in turn solves the differential equation on each of the sub domains where the latest values of the approximate solution of one domain are used as the boundary condition for the other domains.

To illustrate the method, consider Figure 3.1. The top figure represents the complete domain of which the boundary conditions are known and specified. For this example, the (Poisson) differential equation $\partial^2 u/\partial x^2 = -8$ is solved on the domain $x \rightarrow [0,1]$, with boundary conditions $u(0) = 0$, $u(1) = 0$. Of course, the solution is easy to obtain and equals $u(x) = -4\left(x^2 - x\right)$. To illustrate the Schwarz alternating method, the complete domain is now split into three overlapping domains. Domain 1 runs from $x = 0$ to $x = 1/3$, domain 2 from $x = 2/3$ to $x = 1$ and domain 3 from $x = 1/4$ to $x = 3/4$. Consequently, the two overlapping regions range from $x = 1/4$ to $x = 1/3$, and from $x = 2/3$ to $x = 3/4$, respectively. The iteration is started by solving for domain 1 first. One of the boundary conditions is known, $u_1(0) = 0$, but the boundary condition at $x = 1/3$ is not known. However,

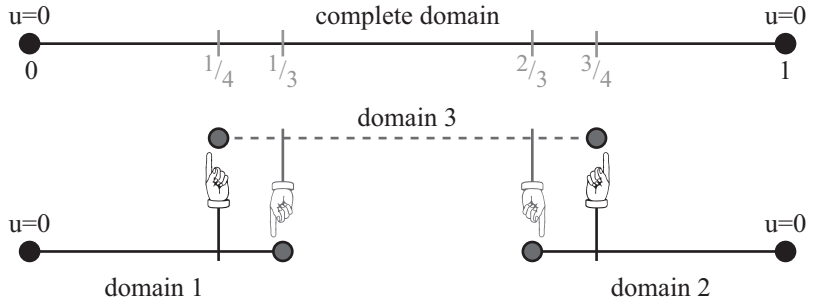

Figure 3.1: Figure illustrating the different domains used in the example. The top illustration shows the complete domain $x \rightarrow [0, 1]$, while the bottom illustration shows how the domain is split into three overlapping domains and how they can communicate with each other.

the Schwarz alternating method states that the approximate solution of domain 3 can be used as a boundary condition for domain 1. Because domain 3 is not yet solved for, an initial guess can be used, which can be "zero" everywhere, making the boundary condition $u_1(1/3) = u_2(1/3) = 0$. This makes the solution of domain 1 $u_1(x) = -4\left(x^2 - \frac{1}{3}x\right)$. The same can now be done for domain 2, where the boundary condition at $x = 1$ is known and the boundary condition at $x = 2/3$ is taken from domain 3 (i.e. the initial guess). The solution for domain 2 is: $u_2(x) = -4\left(x^2 - \frac{5}{3}x + \frac{2}{3}\right)$. At this stage the "new" solutions for domain 1 and 2 are known. These solutions can now be used to obtain a "new" solution for domain 3. Domain 3 needs boundary conditions at $x = 1/4$ and $x = 3/4$. The alternating Schwarz method dictates that the values from domain 1 and 2 can be used. So, from domains 1 and 2, $u_3(1/4) = u_1(1/4) = 1/12$, and $u_3(3/4) = u_2(3/4) = 1/12$, can be used as the boundary conditions for domain 3, making the "new" solution for domain 3: $u_3(x) = -4\left(x^2 - x + 1/6\right)$. The Schwarz iteration can now be continued for domains 1 and 2, where the boundary conditions can be taken from u_3. The second column of Table 3.1 shows how the boundary condition for domain 3, $u_3(1/4)$, changes over the Schwarz iterations, and the third column shows the error relative to the final (exact) solution. The fourth and fifth columns show the same for the value of the solution at $x = 1/2$. So this means that with the current geometry and overlap, the Schwarz alternating method needs about 40 iterations to find the solution for all domains with an error less than 10^{-5} relative to the exact solution.

The example described above gives a good insight in how the Schwarz alternating method works. In this paragraph a closer look is taken on the error analysis and the effect of the size of the overlap used in the method. Figure 3.2 shows these results. The left graph shows how the error from the previous example, decreases with the number of iterations. The error decreases exponentially with the number of iterations. In the previous example, the domain overlap size was picked randomly and kept constant during the iterations. However,

Table 3.1: Convergence of the Schwarz alternating method.

iteration	value $u_3(1/4)$	error $\frac{u_{exact} - u_3(1/4)}{u_{exact}}$	value $u_3(1/2)$	error $\frac{u_{exact} - u_3(1/2)}{u_{exact}}$
1	0.0833	$8.89 \cdot 10^{-1}$	0.3333	$6.67 \cdot 10^{-1}$
2	0.2500	$6.67 \cdot 10^{-1}$	0.5000	$5.00 \cdot 10^{-1}$
3	0.3750	$5.00 \cdot 10^{-1}$	0.6250	$3.75 \cdot 10^{-1}$
4	0.4688	$3.75 \cdot 10^{-1}$	0.7188	$2.81 \cdot 10^{-1}$
5	0.5390	$2.81 \cdot 10^{-1}$	0.7891	$2.11 \cdot 10^{-1}$
10	0.7000	$6.67 \cdot 10^{-2}$	0.9500	$5.01 \cdot 10^{-2}$
20	0.7473	$3.76 \cdot 10^{-3}$	0.9973	$2.82 \cdot 10^{-3}$
30	0.7498	$2.12 \cdot 10^{-4}$	0.9998	$1.59 \cdot 10^{-4}$
40	0.7500	$1.19 \cdot 10^{-5}$	1.0000	$8.94 \cdot 10^{-6}$

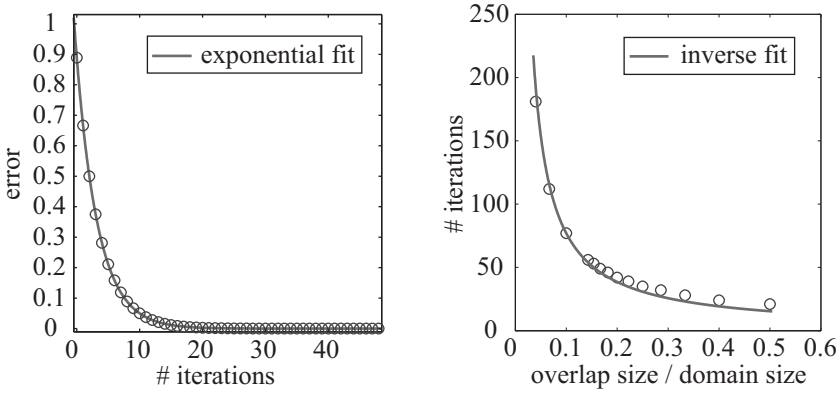

Figure 3.2: Figure showing the error analysis and convergence behaviour of the alternating Schwarz method. The left figure shows how the error for the example shown in Table 3.1 decays exponentially with the number of iterations. The right figure shows the effect of the number of required iterations when changing the overlap size, while keeping the desired minimum error between the exact and iterated solution constant.

the size of the overlap influences the error convergence, or more precise, the ratio between the domain size and the overlap size has an influence on the error convergence. To test this, the same total domain size and function as the previous example is used, although the size of the overlap is varied while keeping the size of domain 3 constant. The right graph in Figure 3.2 shows the result. From the figure it can be seen that the number of iterations that is needed, decreases when the overlap size is increased. The number of iterations in this case corresponds to the number needed to obtain a solution with an error less than 10^{-6}. The datapoints in the figure are fitted to $a/d_{overlap}$, where a needs to be fitted and $d_{overlap}$ is the relative overlap size. This means that the number of iterations scales approximately inversely proportional to the relative overlap size, i.e. $N_{it} \sim 1/d_{overlap}$. Note that when the domain size and overlap size are increased by the same factor, the number of iterations does not change.

So far, only the continuum case was discussed. The Schwarz alternating method can be successfully used to split one large (continuum) domain into several smaller overlapping domains. After a certain number of iterations between the subdomains, the solution for the complete domain is found. This fact and the fact that the results from molecular dynamics simulations produce continuum solutions (provided that the local physics allow it), inspired several researchers [6, 81, 82, 206] to use the Schwarz alternating method to couple molecular dynamics to a continuum, which is illustrated in Figure 3.3. The domain on the left is the continuum domain, where for example the Navier-Stokes equations are solved, while the domain on the right is a molecular dynamics domain, where a MD simulation is performed to solve for the atomic motion.

In order to successfully initiate the iterations between the two domains, two conditions need to be satisfied first. First of all, to start iterating between the two domains, it must be made sure that the boundary conditions on both domains can be specified. As will become clear in the following section, the boundary condition applied to the continuum domain obtained from the MD domain is straightforward. However, the other way around, i.e. setting the appropriate boundary condition for the MD domain, is not so straightforward. The other condition to successfully start the iterations originates from the overlap region. As was mentioned before, an MD simulation gives the right continuum solution, but it will only do so if the conditions for a continuum are satisfied. It is known from previous simulations that, for example near walls, discrete non-continuum effects can be present. When this region is solved with a continuum method, the same result cannot and will not be obtained. Therefore, the position of the overlap region must be placed somewhere where the MD simulation gives the same solution as the continuum method applied in the continuum domain. Note that this position is not always known beforehand and it can become a case of trial and error to find the proper position for the overlapping domain. Finally, the size of the overlap that can

be used in the MD/Continuum case is not yet specified. Obviously, the preferred overlap size is as small as possible, because this will limit the MD domain size, which strongly affects the total computational effort. However, as was discussed above, a smaller overlap size also means that more iterations for the Schwarz alternating method need to be performed before a satisfactory solution is obtained, which increases computational effort again. This means that the ideal overlap size is a compromise between the two. In practice the size of the overlap is just chosen to any convenient value matching either the grid in the continuum domain or the size of the cell list in the MD domain. In Section 3.3 this is further investigated.

As a final note, consider the case what happens if a simulation would be attempted where the overlapping domain is placed somewhere where the continuum domain is not able to solve the same solution as the one obtained by MD, e.g. unexpected constitutive behaviour. In this case the continuum domain, that is unaware of the non-continuum behaviour, will try to enforce this incorrect solution onto the MD domain. On the other hand, The MD domain will continue solving for the non-continuum behaviour, albeit with different boundary conditions. The (end) result will be a incorrect coupled solution and almost certainly diverge, i.e. the convergence behaviour as shown in figure 3.2 will not be obtained. Hence, the position of the overlap domain is crucial, because this is the place where the two domains exchange information. How they exchange information, or how to specify the boundary conditions on the different domains, is discussed next.

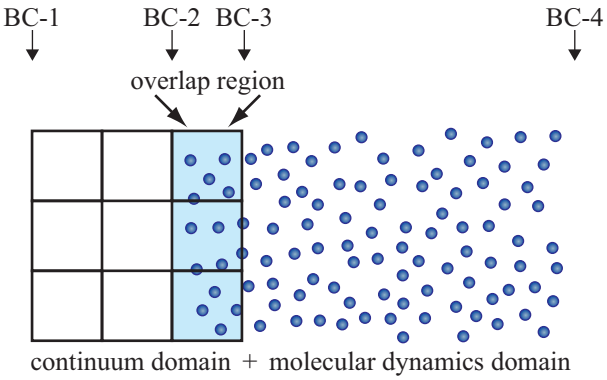

Figure 3.3: The alternating Schwarz method for the MD/CFD coupling. For the quasi one-dimensional case, in total four boundary conditions need to be specified, BC-1 - BC-4.

3.2 Connecting Boundaries (MD→CFD, CFD→MD)

In order for the Schwarz alternating method to work for MD and continuum domains, a specification of boundary conditions on both domains must be available. More specifically, in the quasi one-dimensional coupled case shown in Figure 3.3, four boundary conditions are needed. First, there is the "pure" boundary condition for the far left continuum boundary (BC-1), which could for example be a certain velocity, pressure, or a flux boundary condition. Second, there is a "pure" MD boundary condition for the far right MD boundary (BC-4), which could for example be an atomistic or a specular wall. The other boundary conditions (BC-2 and BC-3) however need further attention. On the right hand side of the overlap region (BC-3), which is the boundary of the continuum domain, a (continuum) boundary condition needs to be specified that is taken from the (approximate) solution of the MD domain at that position. On the opposite side of the overlap region (BC-2), which is the boundary of the MD domain, an (atomistic) boundary condition needs to be specified that conforms to the (approximate) continuum solution at that position. Because of the different variables that are solved in a continuum and in an MD simulation, i.e. continuum variables versus atomistic position and trajectories, clearly these boundary conditions are not straightforward. First in Section 3.2.1 it is explained how the continuum domain can be connected to the MD domain. In Section 3.2.2 the opposite is done, i.e. how the con-

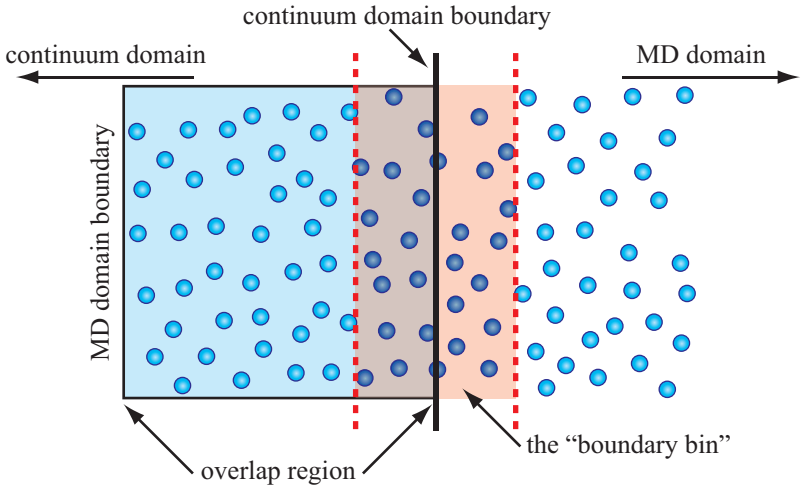

Figure 3.4: Figure demonstrating how to obtain the boundary condition for the continuum (extending to the left) using a bin near the boundary in the MD domain (extending to the right)

tinuum solution can be opposed to the MD domain. Although the techniques explained in this section are general, they do focus on the single atom case, like MD simulation of argon. The coupling of molecules does require some special attention, which is explained in Section 3.2.3.

3.2.1 Connecting the Continuum Domain to the MD Domain

The first issue that needs to be tackled is the connection of the continuum domain to the MD domain. In other words, the issue of the boundary condition that is applied to the continuum domain and is acquired from the (approximate) solution in the MD domain. Luckily, part of this issue was already solved in Section 2.3.2, where the process of binning was explained. There it was established that when a sufficient number of atoms are averaged over a sufficient amount of time, macroscopic variables like density, velocity, pressure and temperature can be obtained from the positions, the velocities, and the interatomic forces of the atoms. Consider the situation sketched in Figure 3.4, where the overlap region is shown. The continuum boundary condition can be acquired from the *boundary-bin*, which is a bin positioned with its center on the continuum boundary and covers a small region near the boundary. Now this boundary-bin can collect data like atom velocities and interatomic forces of all atoms that enter the region during certain time steps. When enough data are gathered to obtain converged statistics, the boundary-bin gives a macroscopic variable, which consequently can be used as the boundary condition. This of course means that the MD simulation must progress for several timesteps before the boundary-bin has collected enough information (see Section 2.3.2 for more details). In other words, the continuum domain does not need frequent updating compared to the MD domain. Please note that it was not specified which continuum variable is binned and coupled. This is because technically any variable in the continuum equation solved in the continuum domain can be coupled this way. This is used in later sections where both velocity and temperature are coupled. Once the boundary condition is known, the continuum domain can be updated with any of the techniques discussed in Section 2.1. This gives a new approximate solution that is applied to the MD domain and eventually will give an improved approximate solution for the MD domain again. However, then first a way of specifying the continuum boundary conditions on the MD domain must be given.

3.2.2 Connecting the MD Domain to the Continuum Domain

The previous section showed how to connect the continuum domain to the MD domain, where it was shown how it is possible to obtain a single continuum variable for the boundary, from many atoms near the boundary at many different time steps. However, in this section the opposite is considered, i.e. knowing a single continuum variable at the boundary, the positions and velocities of the atoms near the boundary at many time steps are desired. Of course, this requirement is unrealistic, since there are many possible sets of atom positions and velocities with same average continuum value. So instead of trying to determine the exact position and velocity of the atoms, a method must be used that at least results in the right average value. This is discussed in Section 3.2.2.1. Then there is another issue; when fluid is flowing into the continuum domain and leaves this domain through the boundary, mass may not be conserved. However, this can easily be solved by applying a flux boundary condition and the continuity equation takes care of the mass conservation by putting constraints on the velocity field. However, the discrete description of MD means that the atoms that actually leave and enter the MD domain, must do so in a way consistent with the flux through the continuum boundary. How this issue is tackled, is explained in Section 3.2.2.2. Then finally, there is an issue that purely exist because of the way MD is computed. Using molecular dynamics, for each atom the interaction with *all* surrounding atoms, or at least all atoms within a certain radius, is considered. On the other hand, the continuum domain, per definition, does not simulate individual atoms. This means that close to the boundary where the MD domain and the continuum domain meet each other, the MD domain still wants to interact with non-simulated atoms inside the continuum domain. If nothing is done, this situation will result in visible non-continuum behaviour near the continuum boundary, which is pure computational and not physical. Clearly this situation must be prevented, which is discussed in Section 3.2.2.3.

3.2.2.1 Imposing the Desired Macroscopic Properties to the MD Domain

As discussed in the previous section, to connect the MD domain to the continuum domain it is required to manipulate the atoms in such a way that they, on average, behave as the continuum. This manipulation (or forcing) is only done in a small region near the MD domain boundary, while the rest of the MD domain is untouched. Figure 3.5 shows the situation that is discussed here. As shown before, the MD domain and continuum domain have an overlap region where both methods should give the same result. In the previous section the boundary-bin was used to communicate the MD solution to the continuum domain. Similar to this, at the MD domain boundary, a *forcing-bin* is now constructed. This forcing-bin is used to communicate the continuum solution to the MD domain. Because of the way the Schwarz alternating method works, this boundary-bin and forcing-bin need to be separated by a certain distance, which as discussed before, influences the convergence of the method (for the results see Section 3.3).

Of course, it depends on the continuum variable that needs to be communicated to the MD domain on how to proceed. However, a straightforward variable to communicate is the macroscopic temperature. In Section 2.3.3 it was discussed how a (full) MD domain could be set to a target temperature using the concept of thermostats. Basically the atoms were manipulated in such a way that the system was solved in a NVT ensemble. Thus it achieves a constant temperature. In the case of the velocity rescaling method, the atom velocities were manipulated in such a way that the kinetic energy, and thus the instant temperature, was constant and set equal to the continuum temperature at every time step. The more advanced Nosé-Hoover method allows for temperature fluctuations while maintaining a constant macroscopic temperature. These methods were previously applied to the whole MD domain, and this can still be done in the coupled MD and continuum domain simulation (including the forcing-bin). However, for the connection of the MD domain to the continuum domain, it is possible to only apply these methods in the forcing-bin, and let the rest of the system respond to the thermostatted area. This now means that every forcing-bin can have its own temperature, which is taken from the continuum domain, and thus enabling temperature coupling as well. The results of such a coupled simulation, where a temperature gradient is applied, is shown in Section 3.3.

Another continuum variable that needs to be communicated is the macroscopic velocity field. In other words, all atoms that are inside the forcing-bin need to be manipulated in

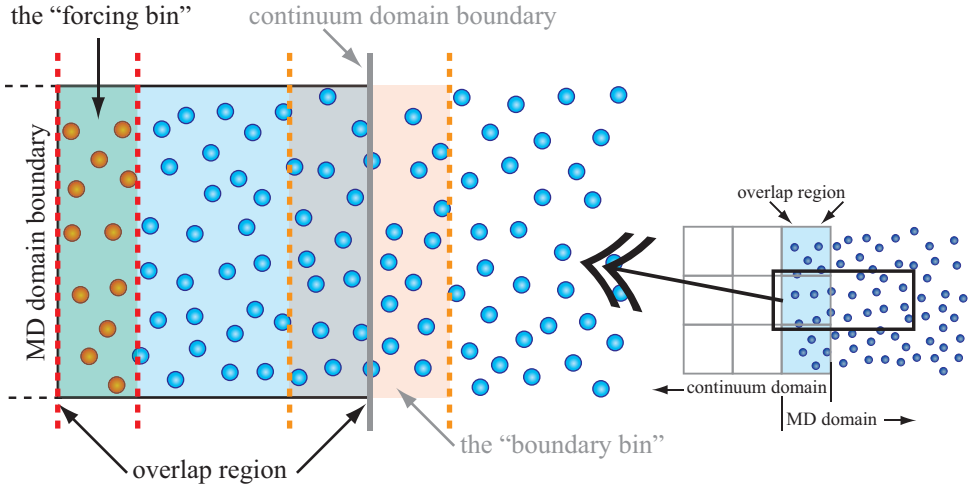

Figure 3.5: Application of the boundary condition for the MD domain obtained from the continuum using a bin near the boundary in the MD domain.

such a way that they on average produce the macroscopic velocity. It needs to be stressed at this point that this does not mean that the velocities of the individual atoms inside the bin are identical, or that the fluctuations in the velocities of individual atoms caused by the finite temperature of the bin vanish. Actually, the method should avoid this all together and can be accomplished as follows. To start the derivation, consider the following equation where the macroscopic velocity is obtained through binning the individual atom velocities and adding through time [40]:

$$\vec{u}_c^t = \frac{1}{N_t} \sum_{n=1}^{N_t} \left[\frac{1}{N_{bin}^n} \sum_{i=1}^{N_{bin}^n} \vec{u}_i^n \right] \tag{3.1}$$

In this equation $\vec{u}_c^t$ is the macroscopic velocity vector at time t, N_t the number of sampled time steps to average the instant velocities, N_{bin}^n the number of atoms that are inside the bin at every time step n, and $\vec{u}_i^n$ is the (instant) velocity vector of atom i at time step n. The velocities of the individual atoms are of course changing at every time step inside the MD simulation, because of the interacting forces between the atoms. The equation expressing the new velocity of atom i is given by (2.57):

$$\vec{u}_i^{n+1} = \vec{u}_i^n + \frac{\Delta t}{2m_i} \left(\vec{f}_i^{n+1} + \vec{f}_i^n \right) \tag{3.2}$$

Here $\vec{u}_i^{n+1}$ is the (instantaneous) velocity of atom i at time step $n+1$, $\vec{f}_i^n$ the total force on atom i at time step n, and $\vec{f}_i^{n+1}$ the total force on atom i at time step $n+1$. These two equations can be combined into:

$$\vec{u}_c^t = \frac{1}{N_t} \sum_{n=0}^{N_t-1} \frac{1}{N_{bin}^n} \left[\sum_{i=1}^{N_{bin}^n} \vec{u}_i^n + \sum_{i=1}^{N_{bin}^n} \frac{\Delta t}{2m_i} \left(\vec{f}_i^{n+1} + \vec{f}_i^n \right) \right]$$

This equation gives the macroscopic velocity in a bin, in terms of MD variables that are binned over time. At this point, the only variable that can be manipulated in this equation in a MD simulation, is the total force on each atom inside the bin at the subsequent time step, $\vec{f}_i^{n+1}$. In Section 2.4 it was shown that with the aid of a body force applied to the whole MD domain, the atoms were set in such a motion that the average motion of the atoms, thus in the macroscopic continuum, was equal to that of a Poiseuille flow. Now, instead of applying the same body force to all atoms, a unique additional force is applied to each atom inside the bin. In other words, the total force on atom i now consists of the interatomic force on i exerted by all surrounding atoms plus this additional force. The value

of this additional force remains unknown for now. So, the total force on atom i is:

$$\vec{f}_i^{n+1} = \sum_{j=1, j\neq i}^{N_{bin}} \vec{f}_{ij}^{n+1} + \vec{f}_{i,extra}^{n+1} = \vec{f}_{i,LJ}^{n+1} + \vec{f}_{i,extra}^{n+1} \tag{3.3}$$

where $\vec{f}_{i,LJ}^{n+1}$ represents the interaction force on atom i, and $\vec{f}_{i,extra}^{n+1}$ is the additional force applied to atom i at time step $n+1$. This equation is substituted into the previous equation, resulting in:

$$\vec{u}_c^t = \frac{1}{N_t} \sum_{n=0}^{N_t-1} \frac{1}{N_{bin}^n} \left[\sum_{i=1}^{N_{bin}^n} \vec{u}_i^n + \sum_{i=1}^{N_{bin}^n} \frac{\Delta t}{2m_i} \left(\vec{f}_{i,LJ}^{n+1} + \vec{f}_{i,extra}^{n+1} + \vec{f}_i^n \right) \right] \tag{3.4}$$

To simplify the derivation at this point, it is convenient to take the mass of all atoms equal to m. This allows to remove the mass from the sum over the bin, and it allows to write the sum of the interaction forces on atom i at time step $n+1$ and the total force on atom i at time step n as average values of the bin. After rewriting the resulting equation for the applied extra force and writing everything as a single sum, the equation becomes:

$$\frac{1}{N_t} \sum_{n=0}^{N_t-1} \left[\frac{1}{N_{bin}^n} \sum_{i=1}^{N_{bin}^n} \left(\vec{f}_{i,extra}^{n+1} \right) \right] = \frac{1}{N_t} \sum_{n=0}^{N_t-1} \left[\frac{2m}{\Delta t} \left(\vec{u}_c - \vec{u}_{avg}^n \right) - \vec{f}_{avg}^n - \vec{f}_{avg,LJ}^{n+1} \right] \tag{3.5}$$

This equation shows how the new value of the additional force on each atom in the bin depends on the force caused by the velocity difference between the (target) macroscopic velocity and the average velocity in the bin at each time step, minus the force that was applied to the bin and the interaction force at the next time step. However, at this point the problem remains that any sequence of values inside the sum for the additional force on atom i can still result in the same (net) value. In other words, the force on each atom inside the bin can still be any value, as long as the LHS and RHS of the equation are equal. An additional assumption is needed in order to proceed. The first assumption is that the extra force is not applied to each individual atom inside the bin, but is constant for the whole bin. The equation for the applied extra force, $\vec{f}_{extra}^{n+1}$, then becomes:

$$\frac{1}{N_t} \sum_{n=0}^{N_t-1} \vec{f}_{extra}^{n+1} = \frac{1}{N_t} \sum_{n=0}^{N_t-1} \left[\frac{2m}{\Delta t} \left(\vec{u}_c - \vec{u}_{avg}^n \right) - \left(\vec{f}_{avg}^n + \vec{f}_{avg,LJ}^{n+1} \right) \right] \tag{3.6}$$

The second assumption is to set the additional force, $\vec{f}_{extra}^{n+1}$ equal to the expression inside the sum on the RHS of the equation. The result of this is an additional force that is applied to the forcing-bin, which can influence the MD solution in such a way that on average the continuum solution is simulated.

$$\vec{f}_{extra}^{n+1} = \frac{2m}{\Delta t} \left(\vec{u}_c - \vec{u}_{avg}^n \right) - \left(\vec{f}_{avg}^n + \vec{f}_{avg,LJ}^{n+1} \right) \tag{3.7}$$

This additional force is a simple way to communicate the continuum solution to the MD domain, and it enables the specification of any (macroscopic) velocity field to an MD domain in the coupled simulation. Note that this expression is similar to the constrained dynamics equation derived by O'Connell et al. [40]. However, at this point, also note that because the additional applied force is set to a single value in the entire forcing-bin each time step, the average value of the velocity *in this bin* is a single constant value. In other words, the averaged velocity profile extracted from the MD simulation with a bin taken as the size of the forcing-bin, will have a plateau over the distance that is equal to the size of the bin. However, if one would sample the velocity profile with a different sized bin, the averaged velocity profile will not be a strict plateau, i.e. absolutely flat, because although the average of *all* atoms inside the forcing-bin is a constant value, this is not the case if only a sub-set of these atoms is sampled.

To demonstrate the plateau or *flatting effect*, figure 3.6 shows the velocity profile obtained from a coupled simulation of Couette flow. Here the MD domain is connected to a continuum domain at two boundaries, while the other boundaries are periodic. Each of the continuum domains have a Dirichlet velocity boundary condition specified on the outside boundary, while the other boundary condition is determined from the MD domain. On the other hand,

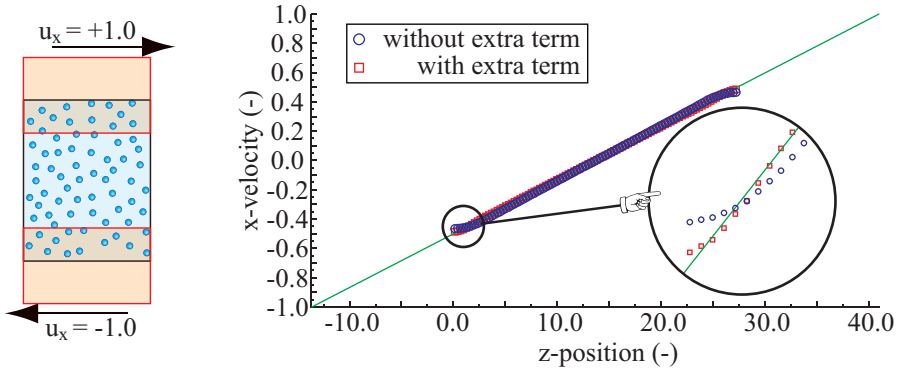

Figure 3.6: Figure demonstrating how the velocity profile does have a plateau caused by the constant extra force inside the forcing-bin and how this can be solved by applying a local shear based on the continuum velocity gradient.

the macroscopic velocity is enforced inside the forcing-bin using the technique discussed above. The resulting velocity profile (dots in figure 3.6) clearly shows a region with a different velocity gradient where the MD domain is connected to the continuum domain, which is caused by applying the same additional force over the entire bin. In order to remove the plateau in the velocity profile, the additional force must be adapted. In the deriviation of the additional force, only the value of the macroscopic velocity was taken into account. In order to remove the plateau inside the velocity profile, also the macroscopic velocity gradients must be included. In other words, inside the forcing-bin also the correct shear needs to be applied. Therefore, starting with equation 2.86, the new additional force in the case of the Couette flow simulated here, can be written as:

$$f^{n+1}_{x,extra,i} = \frac{2m}{\Delta t}\left(u_c - u^n_{x,avg}\right) - \left(f^n_{x,avg} + f^{n+1}_{x,avg,LJ}\right) + \mu\frac{\partial u_x}{\partial z}\Delta z_i \tag{3.8}$$

The difference between this term and the additional force as specified before is that the new term is different for every atom i inside the forcing-bin. The additional term is a shear force determined from the distance atom i is away from the centre of the forcing-bin, $\Delta z_i = z_{centre} - r_i$. The resulting velocity profile (solid line in figure 3.6) is shown in figure 3.6. The plateau is now greatly reduced and the correct velocity gradient is obtained everywhere inside the MD domain. The small difference that remains is caused by another effect that will be discussed in Section 3.2.2.3.

3.2.2.2 Controlling the Mass Flux and Resulting Reinsertion of Atoms

In the previous section it was explained how the macroscopic velocity is communicated to the MD domain using a forcing-bin near the MD domain boundary. However, this combination of macroscopic velocity and domain boundary creates an additional problem. When dealing with a pure continuum flow problem, the boundary conditions are either a specified velocity, a velocity gradient, or a specified pressure at the boundary. The continuity equation dictates the mass conservation of the continuum, and the Navier-Stokes equation dictates the resulting velocity and pressure inside the system. The flow of the continuum through the boundaries can simply be expressed as a mass or volume flux. However, this continuum behaviour now must be translated back to the discrete behaviour in the MD.

The macroscopic variable "mass flux" in the MD simulation can simply be described as follows. Consider an imaginary wall somewhere inside the MD domain; over time, atoms of a certain mass pass this wall from one side to the other, but atoms also do the opposite (i.e. move back through the boundary) because they vibrate due to the finite temperature of the atoms. Each time an atom passes this wall, its mass is added or subtracted from a discrete value representing the total flux, depending on its direction. This motion of atoms through the wall is averaged over time giving the (macroscopic) mass-flux. There are two important things here: 1) atoms move back and forth through the boundary and 2) the average of this movement is the continuum flux through the boundary. However, near the

MD domain boundary this becomes problematic. Once atoms went through the MD domain boundary, they are located outside the MD domain that is simulated, and thus the number of atoms N is no longer constant inside the computational domain. In other words, the (instant) mass of the whole domain changes at that point, and the only way to prevent this is to insert an equal number of atoms into the domain somewhere else. Strictly speaking, this is exactly what the continuity equation imposes in the continuum, i.e. to preserve mass by putting restrictions on the velocity field. However, an atom that went outside the MD domain could be coming back the next time step, and from an averaged point of view, in that case nothing has happened (i.e. no flux). So, in order to communicate the continuum flux to the MD domain, atoms should be able to go through the MD domain boundary, either going out or going in the MD domain. In the first case the atom must not disappear, i.e. somehow remain inside the MD domain volume, and in the second case the atom needs to be reinserted at the right location. Luckily, both problems can be solved, which will be explained next.

First the subject of the insertion of atoms at the right spot is discussed. The first question that arises here is: "What is the right spot?". Inside the MD domain all the atoms interact with each other (and vibrate) through the interatomic potential, which in the case of the Lennard-Jones potential is strongly repulsive at short range and attractive at long range. This also means that those atoms that are in close range with each other have a large (positive valued) energy associated to them, while those atoms that are further away contribute to less (but negatively valued) energy. However, for the NVE ensemble the total energy, the sum of the potential and kinetic energy, remains the same for the whole domain. Therefore, the "right spot" of the newly inserted atom, is a new position inside the domain where the energy of the system does not change too much or remains the same. In other words, the position of the new atom must be determined from the interaction with the atoms that are in the neighbourhood of the new position in such a way that the potential energy associated to the new atom is equal to the presribed (mean) energy. This also means that the insertion of the atom can become computationally very expensive, because for each atom that is inserted many interactions must be determined beforehand. Another problem, especially associated with dense liquids, is the fact that the potential energy landscape can change value of several orders of magnitudes in a very small distance and it has large hills and very deep holes which can act as traps for normal steepest-descent iterators. Delgado-Buscalionia et al. [49] developed a special steepest-descent iterator, called *USHER*, that overcomes this problem by adapting the magnitude of the displacement to the local features of the energy landscape and in this way it can very quickly find an appropriate spot. To prevent trapping, the routine is simply restarted at another random position when it detects it is trapped. The updating of the new atom position $\vec{r}^{n+1}$ is written as:

$$\vec{r}^{n+1} = \vec{r}^{n} + \frac{\vec{f}^{n}}{f^{n}}\delta s^{n} \tag{3.9}$$

where $\vec{f}$ is the (virtual) interatomic force on the atom that needs to be inserted, while f is the modulus of $\vec{f}$ and δs^n is the displacement of the iterator. The USHER insertion algorithm is optimised by choosing the right displacement each time step. The original USHER method suggests an optimal performance by choosing:

$$\delta s^{n} = \begin{cases} \Delta s_{ovlp}, & \text{if } U^{n} > U_{ovlp} \\ \min\left(\Delta s, \frac{U^{n}-U_{0}}{f^{n}}\right), & \text{if } U^{n} < U_{ovlp} \end{cases} \tag{3.10}$$

where U_{ovlp} corresponds to a high value of potential energy ($> 10^4$ in reduced units), associated to the condition where the new atom position is very close to an existing atom, i.e. they almost overlap. The parameter Δs_{ovlp} is chosen such that in the next iteration, the new atom position is moved away from the overlap. This can be accomplished by using the fact that the large repulsive part goes like $U_{repulsive} \sim 4r^{-12}$, therefore chosing $\Delta s_{ovlp} = r_\sigma - (4/U^n)^{1/12}$ will gaurantee that the new atom position is moved a distance r_σ away from the center of the overlapped atom. Here r_σ should be chosen as the characteristic contact distance between the atoms. For the Lennard-Jones potential this is; $r_\sigma = 0.9\sigma$. When the potential energy of the new position is lower than the overlap value, the iterations proceeds in driving the new position downhill in energy, towards the target energy U_0. The two terms between the brackets have different purposes. The second term is used to drive the iterations very quickly into (near) the local minimum, while near the local minimum the

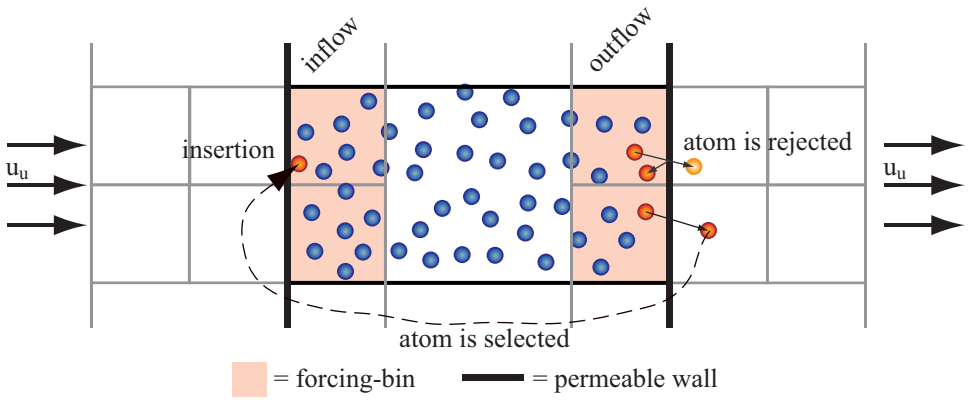

Figure 3.7: The continuum mass flux is applied to the MD domain by means of particle/atom insertion. The atoms that are selected or rejected to be inserted are chosen using a very simple statistical probability.

maximum displacement is governed by the parameter Δs which ideally is chosen near the value of the width of the energy hole.

The USHER iteration is able to quickly find *a* right position with a predefined energy, however, this is only one part of the particle insertion routine. To complete the insertion of the new atom, it must also be given new velocities, but the insertion and possible new position should also take into account the (continuum) mass flux defined by the continuum domain. To illustrate the processes, figure 3.7 shows a simplified view on how this is accomplished. Here the case of uniform flow with a (continuum) velocity, $\vec{u} = \{u_u, 0, 0\}$, is shown. This velocity is enforced with the method described in the previous section in the most left and most right cell. The inflow and outflow of atoms through the MD domain boundaries, where the normal of the boundary is given by $\vec{n}$, must on average be the same as the continuum mass flux $\rho_f (\vec{u} \cdot \vec{n}) = \rho_f u_u$. For the MD domain this means that in order to preserve mass, every time step a certain number of atoms exit the MD domain at the right, while these same atoms are reinserted at the left in the same time step. The average number of atoms that should be inserted every time step Δt is: $n_{insert} = \rho_n A_{MD} u_u \Delta t$, where ρ_n is the number density as defined before and A_{MD} is the outlet/inlet area of the MD domain cell. This would mean that for a typical simulation condition where, $\rho_n = 0.8$, $A_{MD} = 40 \times 20$, $u_u = 0.2$, and $\Delta t = 0.005$, the average number of atoms to be inserted per time step is (only) $n_{insert} \approx 1$. However, if an impermeable wall would be placed at the MD boundary, *all* the atoms that hit the wall are simply bounced back and stay inside the forcing-bin. In the meantime, the enforced velocity u_u inside the cell (which effectively is a force) adds more atoms inside the forcing-bin and under the action of this applied force moves them, on average, towards the MD boundary. This means the density and pressure near the MD boundary is increasing. On the other hand, if no obstruction is placed at the MD boundary, *all* the atoms which go through the MD boundary now are inserted into the cell where there is a inflow (the left boundary). At first, the number of atoms leaving the domain is larger than the number needed based on the mass flux, which means that there will be a depletion of atoms inside the cell. This can mean that now the number of atoms that go through the boundary is less than required by the mass flux, simply because there are not enough atoms available. The result is that the process automatically comes into an equilibrium, where the density inside the cell becomes too low.

From the simplified view in the previous paragraph, it can be concluded that the correct boundary is a permeable wall that allows *some* of the atoms to pass (to be reinserted) while the other atoms are simply bounced back into the outflow cell. This method does require a certain criterion to determine which atoms can pass or not. For example, Werder et al. [206] proposed a criterion based on the collision of atoms with respect to a moving MD domain boundary. This (local) MD boundary is moving with a velocity equal to the macroscopic velocity, but is set back to its original position at each time step. This leaves three options for an atom that was in the vicinity of the boundary. 1) The atom was bounced back by the

moving wall, and has a final position inside the MD domain. 2) The atom bounced back by the moving wall, but has a final position between the position of the wall that was moved and the MD domain. 3) The atom moved outside the MD domain but did not came into contact with the moving wall. The re-insertion criterion is this case is based on the fact that all atoms that are outside the MD domain after the moving wall is moved back, need to be reinserted (i.e. condition 2 and 3).

However, the mass-flux criterion that is used in this work is a simple statistical probability. The first atom that is going to have a new position outside the MD boundary starts with a 50% chance to be selected to be bounced back or to be reinserted. In the first case its velocity component perpendicular to the boundary is reversed while the other components remain the same, while in the second case the atom is removed from the outflow cell to be reinserted into a inflow cell. The selection of this atom is recorded with a certain counter. If the atom was reinserted this counter is increased, while the counter is decreased if the atom bounced back. The number of the counter can now be used to determine the selection for the next atom that is going to have a new position outside the MD boundary. If the counter is positive, this atom has a larger chance to be bounced back, while if the counter is negative the opposite is true. The whole insertion criterion can be controlled by the relative weight that are given to the counter when an atom is bounced back or reinserted. For example, if both events are equally important (+1 for insertion, -1 for bounce-back), after one insertion it is less likely that it is followed by another insertion, but on average both events are occuring with equally chance. However, if different relative weights are given to the events, more atoms are inserted than bounced back (or opposite). Therefore, the relative weigths can be coupled and controlled with continuum information like the density or mass flux (total number of particles inserted) and gives a way to control the mass flux.

The last remaining item for a succesful particle insertion method is to find a way to apply the correct mass flux in the cells that have an inflow condition. In the simplified situation illustrated in figure 3.7 this was very simple; every atom that left the domain at the right had to be reinserted into the cell at the left. However, in general, multiple cells have an inflow and an outflow condition and once an atom has left one of the cells, it must be inserted in *one of* the inflow cells. Similar to what was the case with the outflow mass-flux, the velocity that is enforced inside the inflow cell will move the atoms, on average, towards the center of the MD domain. When there are not enough atoms inserted (or provided) inside this cell, the cell will be depleted from its atoms on one side. This situation is worse when the mass flux inside the cell is higher. Based on this reasoning, the cells should receive the amount of atoms based on their relative mass-flux with all cells. In other words, the cell with the highest mass-flux will have more chance of receiving a new atom.

This ends this section where it was demonstrated how the continuum mass flux and resulting particle insertion can be accomplished. This method, together with the method discussed in the previous section, is all which is needed to succesfully couple the MD domain to the continuum domain for steady incompressible flow. However, the fact that the MD domain has a non-periodic boundary condition where it is connected to the continuum domain and the fact that in molecular dynamics one atom interacts with *all* its neighbours brings forth a problem that manifest itself in unrealistic non-continuum effects. These will be discussed in the next section.

3.2.2.3 Prevention of Fluctuations Near the Boundaries

Molecular dynamics works on the principle that the force on one atom is obtained by looking at the interatomic interaction of all surrounding atoms, or in the case of a cutoff radius, only those inside the radius of the one atom. Whenever a "pure" MD domain is modelled, this is not a problem, because even at the boundary of the MD domain, the atom interacts with the atoms on the other side of the MD domain because of the periodic boundary condition. The result is that this one atom cannot make a distinction between a bulk condition and the post-processing results show a constant density profile, and other profiles. On the other hand, when the interatomic interaction of an atom near a molecular wall is examined the density profile will show variations. These are the (real) wall-fluid interactions that are explicitly modelled. However, with the coupling of the MD domain and the continuum domain a different problem arises. At some point the MD domain stops and the continuum domain is the only computational domain left. Near this boundary is where the particles get inserted and may leave the domain, possibly to be reinserted somewhere else. However,

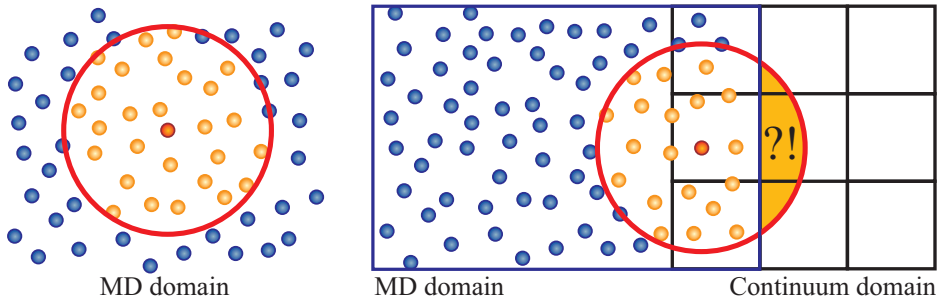

Figure 3.8: Figure demonstrating how the atoms near the boundary have incorrect interatomic interaction

the atoms inside this area do have incomplete interactions, and this result in unrealistic variations in the density profile and other profiles. The missing interaction is illustrated in Figure 3.8 by the shaded area inside the radius of interaction.

The incomplete interaction is a results of the fact that the continuum domain does not simulate atoms or molecules by definition, while molecular dynamics *requires* the individual position of the atoms in order to compute the interaction. This also means that the amount of missing interactions differ as a function of the distance to the boundary and that by providing the molecular dynamics simulation some meaningful information, the missing information can be compensated. For example an atom that is relatively far away from the boundary will only miss a small part of the long-range attractive tail of the Lennard-Jones potential, while an atom that is very near the boundary at the right, predominately interacts with atoms on its left side, which translates back to a strong force towards the right, which in both cases is easy to compensate in order to complete the total interaction. However, the missing interaction potential and force on an atom that is in between the two extreme cases is a complicated function that mostly depends on the structure of the liquid which in turn depends on variables like density, pressure, temperature, and their gradients. The fact that the local structure of the liquid plays a role in the function for the missing interaction was also proposed by Werder et al. [206] where the radial distribution function $g(r)$ describing the local structure for a monatomic fluid like argon, is used. The model for the missing interaction is based on the integration of the force components normal to the boundary and the potential energy contributions weighted by $g(r)$ over the shaded area in figure 3.8. They also showed a way how the missing interaction can be determined from a periodic MD simulation. As explained above, the missing interaction is due to the fact a certain amount of atoms are not taken into account, where the amount depends on the distance from the boundary. The idea of the periodic MD simulation is to simulate a certain atom inside a MD domain that has a complete set of interactions all the time. However, at the same time the total instant force on the atom is split into one part that originates from all those atoms at one side of a virtual wall, while the second part contributes to the interaction of the atoms on the other side of this virtual wall. This last contribution is the part that an atom near the MD/continuum boundary would have missed. If this is done for several virtual walls and sampled for several time steps, this method should give the missing interaction as a function of the distance from the MD/continuum boundary.

The solid line in the left graph in figure 3.9 shows this function for a MD simulation where the number density $\rho_n = 0.8$ and the reduced temperature is $T = 1.2$. As predicted by the extreme cases, the missing interaction force has a weak-attractive force at long distance from the MD/continuum boundary, while the force is large-repulsive whenever the atom is very near the boundary. Please note that the curve between those two cases shows a dip where the force is strong-attractive. The right graph in figure 3.9 shows how the normalised density profile ($\rho^* = \rho/\rho_n$) near the MD/continuum boundary changes when this boundary force function is used. The circles represent the normalised density profile when no boundary force function is used. In this case the density profile shows large variations near the boundary, indicating layering of atoms. This layering is the result of atoms being bounced back at the boundary while they at the same time only have interactions with atoms on one side, thus forcing them towards the wall on average. However, because now there are more atoms in the

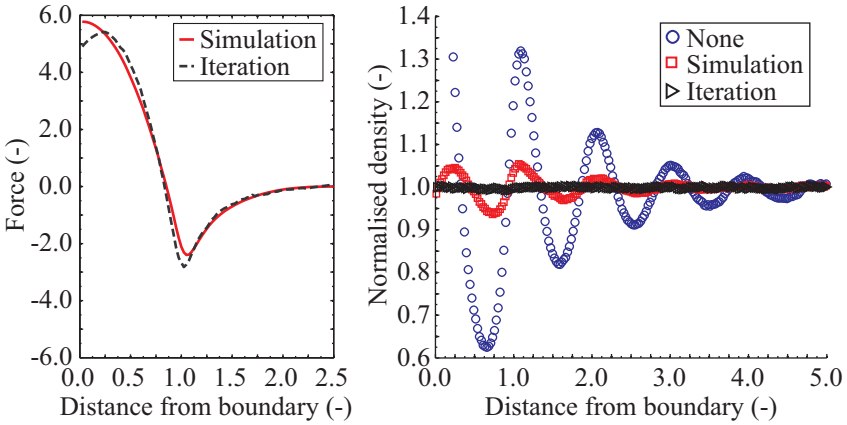

Figure 3.9: Figure demonstrating the density profiles (right) obtained using different boundary functions (left). The boundary functions show how the missing interaction is summarised in one function that is dependent on the distance from the MD/continuum boundary.

region very near the boundary, the large repulsive force at short distance of the atoms result in a dip in the density profile further away from the boundary and both effects are repeated a couple of times away from the wall with lowering amplitude. The density profile obtained when the boundary force function is used, is shown with the red squares and gives a much better result. Although the variations in the density profile are much less, some variations do remain. Very near the boundary the normalised density is less than 1, indicating that the force in this region is too large. The remainder of the profile show a similar pattern as the case when no boundary force is used. The amplitude of the variations is considerably lower, however the peaks and dips of this profile are slightly shifted.

The fact that some variations remain, means the boundary force is still missing some interactions or the conditions near the MD/CFD boundary are not completely the same as those in the bulk condition. One fact could be the influence of the bounce-back wall that regulates the mass-flux, while others [110] think the lack of fluctuation in the boundary force are the cause of the differences. However, as shown above, by adding the boundary force, the density variations became less, which also means that it is possible to change the boundary force in such a way that there are no density variations anymore. It was also already shown qualitatively that the boundary force near the boundary is too high, while the rest of the boundary force must push more atoms into the low-density regions, while at the same time pull them from the high-density regions. In order to quantify this new boundary force function an iterations with respect to the resulting density needs to be performed. Kotsalis et al. [110] showed how this can be done using a very simple control algorithm. In this algortihm the boundary force is slightly changed every iteration, depending on the sampled density profile obtained from the coupled MD/continuum simulation. The change in boundary force at a certain point of the boundary force function is computed using the local gradient of the density profile. This basically is the push and pull analogy explained above, where a dip in the density profile means that atoms on average are pushed into this dip from both sides of the dip. Whenever the density profile is very near the wanted value, the resulting change in force is near zero. The dotted line in the left graph of figure 3.9 show the iterated function for a MD simulation where the number density $\rho_n = 0.8$ and the reduced temperature is $T = 1.2$. Compared to the function obtained with the periodic simulation, the largest changes are near the MD/continuum boundary where indeed the value is lower, while the dip in the boundary force function is slightly shifted and is larger. The black triangles in the right graph in figure 3.9 shows how the normalised density profile is everywhere near one, meaning that the iterated boundary force eliminated the variations in density near the boundary.

Although the boundary function obtained from the iteration with respect to the density profile is a correct boundary force function that completely removed the density variations near the MD/continuum boundary, its shape raises a question. The solid line in figure 3.10

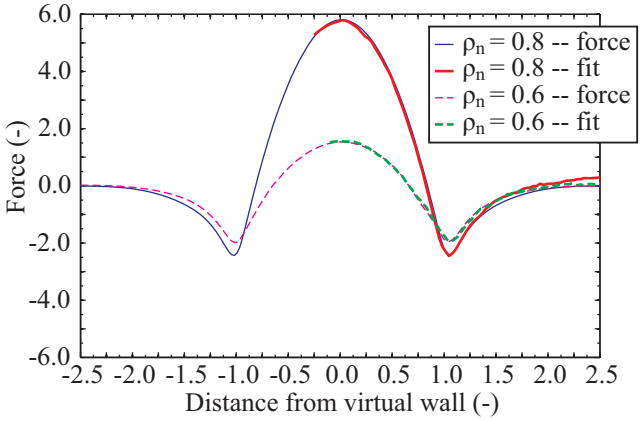

Figure 3.10: Figure demonstrating the average force on one atom as function of the distance from the atom obtained from a periodic simulation. The figure also shows how part of the function, after simple scaling or stretching, is very similar to the boundary force function obtained by iteration with respect to the density profile.

shows again the boundary force extracted from the simulation of the periodic MD domain with $\rho_n = 0.8$ as done before, however this time the function is also extracted for the other side of the atom. For example, the negative numbers in the distance from the atom indicate left of the atom, while positive numbers indicate right of the same atom. As expected, this function is symmetric with respect to the position of the atom ($r_d = 0$), but also the derivative of this function at $r_d = 0$ is zero. However, when these results are compared with the boundary force function that completely eliminated the density variations (shown as the dotted line in the left figure of figure 3.9), it is evident that when this function would also be symmetric at $r_d = 0$, the derivative of the force would not be zero at $r_d = 0$. Another observation of the differences between the two functions showed how the first dip in the curve is shifted. On the other hand, the derivative *is* zero at the top of the first peak at $r_d \approx 0.25$, which again indicates a shift. Comparing these two shifts, reveals that the new boundary force function can be made to match the solid line in figure 3.10 by shifting and stretching it a little bit. The result of this is shown in figure 3.10, where this is done for a simulation with $\rho_n = 0.8$ and $\rho_n = 0.6$. The conclusion from this comparison is that the simulated and iterated functions are very similar in shape, however to completely eliminate the density variations, the simulated function needs to be shifted and stretched. Further research should point out what the cause of this shift and stretch are.

This ends this section where it is explained how large density variations appear near the MD/continuum boundary. These are caused by the missing interactions from the atoms that by definition of a continuum are not modelled. However, it was shown how the missing (averaged) interactions can be supplied by an extra boundary force depending on the distance from the wall and how this function can lower or even elimate the density variations. This boundary force function can either be obtained using a separate periodic MD simulation or by iteration together with the observed density profile.

3.2.3 Connecting Boundaries when Modelling Water

In the previous section it was shown how the MD domain and the continuum domain were connected using the Schwarz alternating method. The coupling of the two domains relies on the fact that the boundary conditions of each domain could be applied to, and extracted from each subdomain. The theories explained so far were relatively general, but concentrated mainly on the case where there is only one type of *atom*, i.e. argon. In this section it is shown how comparatively the same techniques can be used to couple *molecules*, especially water.

In the case of molecules, the connection of the continuum domain to the MD domain remains

the same. The boundary conditions for the continuum equation can be extracted by binning or averaging the results obtained from the MD domain. Although the simulation of the molecules in the MD domain is much more complex, the difference in the continuum domain is simply specified by using different continuum parameters like density and viscosity of the liquid that is simulated.

However, the connection of the MD domain to the continuum domain needs more attention. As some in the previous subsection, there are three vital parts for a successful coupling of the MD domain to the continuum domain. The first part is: 1) opposing the correct macroscopic variables, e.g. velocity or temperature, to the MD domain. In the Argon case, this was done by an additional force inside the forcing-bin to oppose the correct macroscopic velocity, while the macroscopic temperature could be opposed by using a local thermostat. The second vital part is: 2) generating the macroscopic mass flux inside the MD domain. This involves a controlled method that selects which atom does exit and re-enter the MD domain. Whenever an atom is selected to be reinserted, this atom must be inserted in a spot that has the right contribution to the total energy of the MD domain. The third vital part is: 3) to minimise or eliminate unnatural variations in obtained results from the MD domain. These variations are the result of incomplete interactions of atoms that are near the MD/continuum boundary, because there are no atoms simulated inside the continuum domain. This part is solved by adding a mean boundary force to those atoms that are near the boundary. The value of this force is dependent on the distance between the atom and the boundary.

For general molecules, e.g. water, opposing the correct macroscopic variables, the method explained in Section 3.2.2.1 is still applicable. The only difference is that in the case of general molecules, the value of the mass of atom i, m_i, cannot be assumed as one value. For Argon this was done in order to derive equation (3.6), which is the equation that gives the additional force inside the forcing-bin. However, a similar equation can be derived in the case of general molecules by considering the averaged mass inside the cell. In other words, the additional force in the case of general molecules becomes:

$$\vec{f}_{extra}^{n+1} = \frac{2m_{avg}^{n}}{\Delta t}\left(\vec{u}_c - \vec{u}_{avg}^{n}\right) - \left(\vec{f}_{avg}^{n} + \vec{f}_{avg,LJ}^{n+1}\right) \tag{3.11}$$

where m_{avg}^{n} is the average mass inside the forcing-bin at time step n. Consequently, the equation of the additional forces that includes the velocity gradient, 3.8, is changed accordingly. The method to opposing the macroscopic temperature to the MD domain remains the same, i.e. using a local thermostat.

The methods used to oppose the macroscopic mass-flux to the MD domain for general molecules is very similar to the methods used for Argon. However, the big difference is the fact that whole molecules (consisting of bonded atoms) must be able to exit and enter the MD domain and can rotate, which generates several problems. It now is possible that certain atoms (belonging to a molecule) can be *outside* the MD domain but should still be included in the interatomic interaction evaluation. Even now atoms are outside the MD domain, the molecule itself does not necessarily need to be reinserted because of this. It is only when the molecule's center of mass crosses the boundary, the selection method, whether the molecule needs to be reinserted or bounced-back into the domain, is invoked. If it is bounced-back, the whole molecule (all separate atoms) is bounced-back. If it is selected for reinsertion, the whole molecule (all its atoms) is removed from the domain to be reinserted somewhere else. The reinsertion of one Argon atom was a relatively easy task because only one atom and its interaction with all surrounding (single) atoms needed to be considered. Special techniques, like USHER [49], can speed-up this process of finding the right spot inside the energy landscape of the MD domain, based on the instant total energy of the system. However, the reinsertion of a whole molecule is a formidable task, because the interaction of *all atoms inside* the molecule with all surrounding atoms need to be computed. Further, in the case of a single atom, like Argon, there was no need to also find an orientation of the atom. In other words, in the case of Argon only a new position was needed. In the case of a general molecule, the positions of all atoms inside the molecule also result in a certain orientation of the molecule. In the case of fixed bonds and angles between the atoms inside the molecule, this means for example that the new center of mass and the new rotation of the reinserted molecule needs to be computed. Luckily, for the relatively simple molecules (in terms of configuration), especially water, there are techniques available in literature [60]. Similar to the technique used in USHER [49], De Fabritiis et al. [60] developed a more

general USHER technique that can be used for polar liquids, where the energy minimisation process, described in Section 3.2.2.2, is now applied concurrently to all degrees of freedom (translational and rotational).

The last minimum requirement for successful coupling of the MD domain and the continuum domain is the elimination of the unnatural variations in the MD results near the MD/continuum boundary. In the case of an argon atom, these unatural variations were contributed to the missing interatomic interactions due to the non-existing (virtual) atoms inside the continuum domain. In Section 3.2.2.3 it was explained how these missing interactions could be simulated by a (single) averaged boundary force function, where the value of the missing interaction is dependent on the distance from the atom and the boundary. This function can be obtained from a separate simulation of a periodic domain where all the interactions on a single atom are collected as a function of the distance between the atom and a virtual (non-existing) wall. Although this obtained function is successful in reducing the variations, the variations only dissappeared completely when the boundary force function was adjusted slightly. This adjustment was accomplished by an iteration where the desired density profile and the obtained density profile were compared and where the boundary force was changed accordingly (see Section 3.2.2.3 for details). Of course, the same technique can be applied to any molecule. If the correct force is applied to the molecule, then from a macroscopic point of view, the correct density can be obtained or enforced by the iteration. This can be done because the density is simply defined as a certain amount of atoms (mass) that are present in a certain volume (a bin), *on average*. However, the same problem that was encountered with the reinsertion of the molecules also occurs here, namely the orientation of the molecule. Additionally, another problem must be solved. To be precise, the inclusion of the long-range interactions, like Coulombic interactions.

One attempt to percieve such a boundary force for the use of molecular dynamics simulations of water, was done recently by Kotsalis et al. [109] and further applied very recently by the same author [108] to a lipid bilayer in water. Here the density/boundary force iteration method the same author developed before [110], is applied to water. The iteration is performed between the obtained and desired density profile to find a single boundary force that can be applied to the centre of mass of each water molecule. Further, this simulation used a fixed cutoff radius for both the Lennard-Jones interactions and the Coulombic interactions of about 1 nanometre, effectively not simulating the long-range interactions. Using this technique, the density variations in the density profile of water can completely be eliminated. However, it is also shown [109] how the probability distribution of the orientation of the water molecules near the MD/continuum boundary is very much different than the distribution somewhere else in the domain. In other words, if nothing is done to prevent so, the water molecules have a certain preferred orientation near the boundary. This preferred orientation can be the result of other missing interactions (e.g. long-range interactions), the influence of the specular wall, or simply because the boundary force is applied to the centre of mass. To investigate the influence of the position where the boundary force is applied, an adaption of the density/boundary force iteration method is developed here. The difference between this method and the method described in [109] is the fact that now for each atom inside the water molecule, a separate boundary force is obtained and applied at the position of the atom itself (instead of the centre of mass of the molecule). Of special interest is the question whether this has any effect on the preferred orientation of the water molecule near the MD/continuum boundary. The detailed results of this simulation can be found in Section 3.3.4.

In conclusion, the same techniques that are used to couple a MD domain containing argon atoms and a continuum domain, can be used to couple MD domains that contain molecules. Especially in the case of small molecules, like water, these techniques can easily be adapted. However, it must be appreciated that for larger molecules several issues, e.g. orientation of the molecule, still remain and that the computational effort can become a limiting factor.

3.3 Results

In this section several applications and results of coupled continuum and MD simulations are given. In Section 3.3.1 it is demonstrated that the convergence of a coupled MD and continuum simulation is very similar to that of the pure continuum method when both

are used to model continuum bulk behaviour. In Section 3.3.2 it is shown how a coupled simulation using argon can be used to model Poiseuille flow inside a nano channel. The non-continuum effects near the wall are modelled with MD, while the remainder of the domain can safely be simulated using any continuum method. In Section 3.3.3 the coupling of an MD domain and a continuum domain is extended to two-dimensional coupling. Here the flow around an MD obstacle is shown. This effectively also means that a MD simulation is performed where only continuum boundary conditions are applied to the system, i.e. no periodic boundary conditions are needed. In Section 3.3.4 a coupled simulation using water is shown. First the results of the boundary force function, as discussed in Section 3.2.3, are shown. Using these results, Poiseuille flow inside a nano channel is simulated, where the water molecules near the boundary are simulated using MD, while the remainder of the channel is modelled using a continuum method. Finally, in Section 3.3.5 a coupled simulations is shown, where the coupled variable is different that the (macroscopic) velocity. The simulation shows how a temperature gradient applied to the continuum domain can influence an atomistic particle inside the MD domain. It is also shown how the same principle can be applied to a charged particle and its surrounding ions inside the MD domain and how it can influence the continuum domain.

3.3.1 Convergence Tests for the Schwarz Alternating Method

In Section 3.1 the Schwarz alternating method was introduced. There, certain convergence and error properties were identified when the method is applied to continuum domains only, i.e. only partial differential equations. It was explained that this method could also be used to couple a MD domain and a continuum domain, as long as certain conditions are fulfilled. In this section, for that case, a validation is given for the behaviour of the error with the number of iterations, and how the required number of iterations changes with the size of the overlap region.

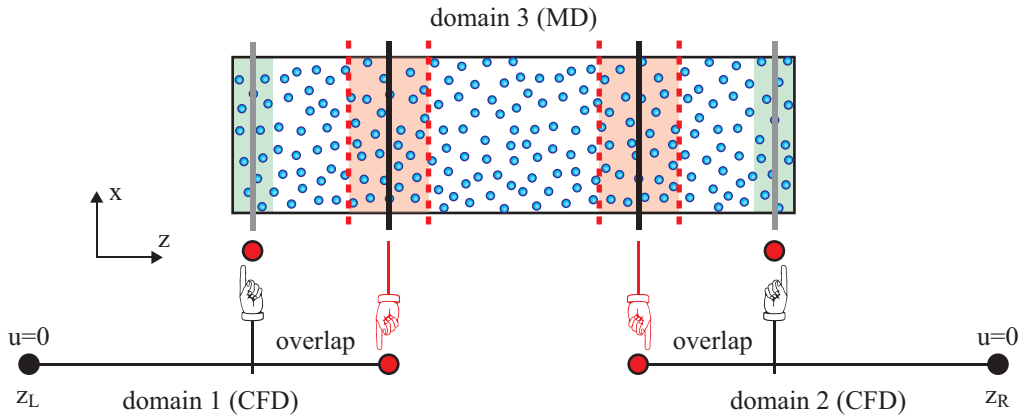

Figure 3.11: The model layout used for the convergence test. Domain 1 and 2 are continuum domains, while domain 3 is simulated with MD. The applied boundary conditions are indicated in the figure. The boundary conditions for the MD domain that are not relevant to the coupling are chosen to be periodic.

To test the convergence of the error, consider the model layout shown in Figure 3.11. Here three domains are specified exactly as was done in the example in Section 3.1, where all three domains were solved as a continuum, and all three were solved for the same differential equation. However, here only domains 1 and 2 are continuum domains, while the domain in the middle, domain 3, is an MD domain. As before, the Schwarz alternating method is initiated with a guessed initial solution, while the two outer boundary conditions for domains 1 and 2 are set to a certain value. However, the difference now is that the behaviour in domain 3, i.e. the MD domain, is not described as a single differential equation with boundary conditions. This means the discrete behaviour of domain 3 somehow must be manipulated in such a way that basically the same differential equation is solved for.

One of such manipulations can be the solving for a Poiseuille flow in all domains, where the velocity is only in the z-direction, while the coupling is done in the x-direction (see Figure 3.11). The continuum differential equation to be solved is the stationary Stokes equation 2.6 with an externally applied force:

$$\mu \frac{\partial^2 u_x}{\partial z^2} = -\rho f_{x,ext} \tag{3.12}$$

In this equation, μ is the fluid viscosity, u_x the velocity in the x-direction, ρ the fluid density, and $f_{x,ext}$ the externally applied force in the x-direction. The solution of this equation is of course a parabolic velocity profile. However, Poiseuille flow was also simulated in a (pure) MD domain in Section 2.4, where the flow was obtained by applying a body force to each atom. Together with the simulated MD walls that created a nano-channel, the resulting velocity profile was parabolic. However, these same walls also caused non-continuum flow near the boundary, caused by the wall-fluid interaction. On the other hand, at the centre of the nano-channel the flow was exactly the same as that of a continuum flow. Therefore, in this Schwarz convergence test, the MD domain is only used for that part of the velocity profile where a continuum was simulated. Poiseuille flow can be obtained by applying a body force to each atom. This simulation method ensures that in all domains the same conditions are simulated, e.g. continuum flow, by different simulation techniques. Initial solutions for the MD domains were discussed in Section 2.3.2, where it was shown that the position of atoms were placed in according with a certain lattice (e.g. FCC) and the velocities of the atoms obey Maxwell's distribution at a given temperature. The system is then "melted" into a liquid form. However, for the Schwarz alternating method, generally an initial solution of zero velocity everywhere in the domain is used. This is a simple concept for the continuum domain, but this means that for the MD domain the system first must melt, while no additional force is applied. Once the system has molten and the binned velocity of the MD system equates to zero, the Schwarz alternating method can be initiated. In practice this is not a stringent condition, because generally the MD system equilibrates faster than it takes to collect enough statistics for one Schwarz iteration update.

Now that the domains and initial conditions are specified, the boundary conditions need to be specified. For the two continuum domains, the boundary conditions are very straightforward. To simulate Poiseuille flow in a continuum, the two outer boundaries simulate a no-slip wall. So, the left boundary of domain 1 and the right boundary of domain 2 have a boundary condition $u_z = 0$. All other boundary conditions must be obtained by communication between the MD domain and the continuum domains as discussed in this chapter. In a simulation the continuum domains are updated every 250 time units, which is equal to 62500 MD time steps in this case. The total MD simulation time is 6×10^6 time steps. This means that in total 96 Schwarz iterations are carried out during the complete simulation. After each iteration the MD solution is used to update the continuum solution, and during the next iteration the new continuum solution is used to update the MD solution. It is expected that the number of Schwarz iterations is more than enough for a converged solution. However, the simulation needs to be extended in order to obtain meaningful statistical information for the average velocity profile from the MD simulation when the coupled solution has actually converged.

The total number of atoms that are simulated is $N = 4608$, while the density of the liquid is $\rho = 0.8$ and the temperature is kept constant by a Nosé-Hoover thermostat at $T = 1.2$. The value for the applied external force is specified as $f_{x,ext} = 0.00125$. The only remaining variable to be specified is the viscosity. Unlike in the continuum, the viscosity in a MD simulation is a result rather than a specified variable (see Section 2.4). In an MD simulation the substance modelled has a certain viscosity based on conditions like the temperature of the MD domain. The MD simulation here simulates argon and in Section 2.4 it was already shown that the viscosity of argon at $T = 1.2$ is $\mu \approx 2$. So, the viscosity for the differential equation that describes the motion of the continuum is specified on that basis, $\mu = 2.0$ units. This makes the solution of the differential equation for the velocity as follows:

$$u_x(x) = \frac{\rho f_{x,ext}}{2\mu}(z - z_L)(z - z_R) = -0.00025(z - z_L)(z - z_R) \tag{3.13}$$

where z_L is the left boundary of domain 1 and z_R is the right boundary of domain 2, which values are given next.

The last prerequisite for the Schwarz iterations is the size of the overlap. As explained in

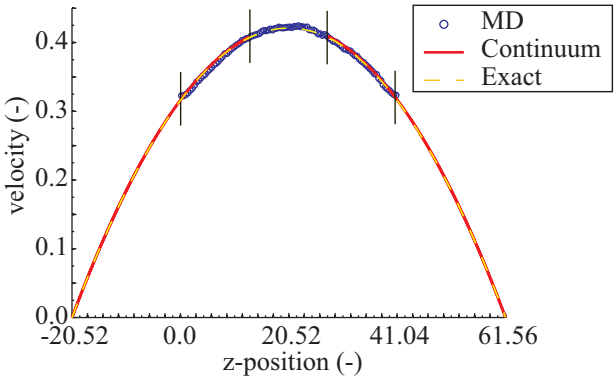

Figure 3.12: The velocity profile obtained in the coupled simulation when an overlap of 8 cells is used. The vertical lines indicate the boundaries of the coupled domains.

Section 3.1, from a computational point of view, the choice for the size of the overlap is governed by the layout and the total number of cells used in the cell list method. For this simulation, the MD domain has a size of $41.04 \times 10.26 \times 13.68$, and it is subdivided into 28 cells in the x-direction. Here the two top and two bottom cells are used to communicate the continuum solution to the MD domain (i.e. the forcing-bin). The boundary-bin, which is used to communicate the MD solution to the continuum domain, also has the size of two cells, which improves the rate of convergence of the statistics. The overlap used for the Schwarz alternating method is determined by the distance of the centres of the bins. For example, when the boundary-bin is placed at cell numbers 5 and 6, the overlap runs from the top of the first cell up to the top of the fifth cell, making the overlap a total of 4.0 cells. For the results it will however be more convenient to specify the overlap relative to the size of the whole MD domain. To check the dependence of the convergence on the size of the overlap, five different overlaps are simulated. The overlaps that were used are 2, 4, 6, 8, and 10 cells. Finally, the left and right boundaries of the continuum domains need to be specified. The size of the whole coupled domain (continuum and MD) is fixed to twice the size of the MD domain, and the MD domain is placed exactly in the middle. This means that $z_L = -20.52$ and $z_R = 61.56$.

Next the results are discussed. As expected, the coupled solution is sufficiently converged (with a relative error less than 10^{-6}) in all cases with a different overlap, before the completion of the 96th Schwarz iteration. Therefore, a converged velocity profile can be obtained using the last 10^6 MD time steps of the coupled simulation. Figure 3.12 shows the velocity profiles in both the continuum domains, the velocity profile obtained from binning in the MD domain, and the exact solution of the whole domain, for the simulation with an overlap of 8 cells. The vertical lines in the figure indicate where the boundaries for the coupled domains are located. The results show that the simulation produces a Poiseuille flow, and that the Schwarz alternating method matches the exact solution in the overlap region, and that both the MD domain and the two continuum domains responded accordingly. The small visible difference (less than 1%) between the exact solution and the coupled solutions can be explained by two predominant causes. The first one being the statistical noise in obtaining the averaged (macroscopic) velocity from the MD domain, and the second one from the fact that, as explained before, the viscosity of the MD domain is a result of the simulation and not a predefined parameter as is the case in the continuum. A small difference between the two values of viscosity can result in small deviations in the final solution.

Now it is established that indeed Poiseuille flow is obtained and that all domains give the same solution, it is checked how the convergence of this solution behaves. The left graph in Figure 3.13 shows how the error relative to the exact Poiseuille solution decreases as a function of the number of iterations for the case where the overlap size is 8 cells. Each point in this curve represents the 62,500 MD timesteps that needs to be simulated to obtain a new continuum solution. The fitted curve in the figure is exponential. It is evident that, just as shown in Section 3.1 with the pure continuum case, the error in the coupled MD

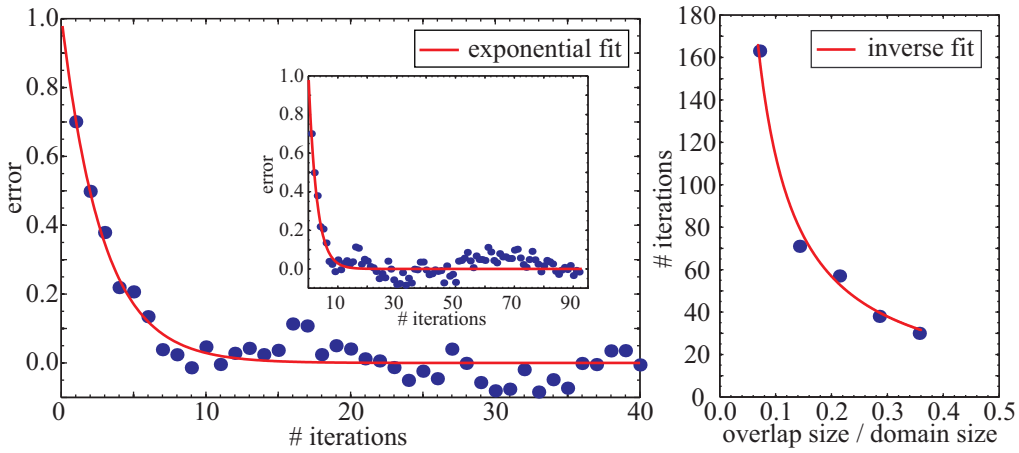

Figure 3.13: The convergence of the error for Schwarz alternating method and the number of iterations as a function of the size of the overlap (for a relative error less than 10^{-4}). The solid line indicates the fitted curve; the error is fitted with an exponential function, while the number of iterations is fitted with an inverse function.

and continuum solution decreases exponentially. However, it must be noted that, especially when the error becomes very small, the statistical noise in the MD solution becomes obvious. This statistical noise can only be lowered when the number of atoms in the MD simulation is increased, or when more time steps are used (i.e. more statistics are gathered). This means that the fitted curve is used, in order to determine the number of iterations that it takes to achieve a solution within a given error. Doing this for the case shown in the figure, the coupled MD and continuum solution has a relative error less than 10^{-6} when 38 Schwarz iterations are performed. The number of iterations that are needed to obtain a relative error less than 10^{-6} for different sizes of the domain overlap are shown in the right graph of Figure 3.13. The points in the figure are fitted to a curve where the number of iterations is inversely proportional to the overlap size relative to the domain size. Again, as in the continuum case, the number of iterations scales inversely proportional with the relative overlap size.

This section showed that the convergence of the error for the Schwarz alternating method in the case of a coupled MD and continuum simulation behaves exactly the same as in the pure continuum case. It also showed that by using this method, the MD and continuum domains can communicate with each other, and when the solution has converged, both the MD and continuum domains yield the same solution.

3.3.2 Comparison of a Pure MD and a MD/CFD Coupled Model

In the previous section it was shown how the Schwarz alternating method can be used to couple two continuum domains and one MD domain to simulate continuum Poiseuille flow in a channel. The continuum domains were applied near the wall, while the MD was applied in the middle of the channel. This was to ensure that all domains solve a continuum solution. In this section the opposite is done. In Section 2.4 it was shown that Poiseuille-like flow can be simulated by a pure MD simulation. The conclusion of that simulation was that in the middle of the channel a good matching Poiseuille velocity profile could be obtained. However, near the walls the wall-fluid interaction causes non-continuum behaviour, like density variations. Therefore, in this section a coupled simulation is performed where the MD domain is used to simulate the non-continuum behaviour near the walls, while the rest of the channel is simulated as a continuum domain. To verify and compare the results, a full MD simulation of the channel is also performed.

Figure 3.14 shows the model layout used for the coupled simulation. The fluid near the wall, including the wall itself, is represented with atoms and is solved with MD. The con-

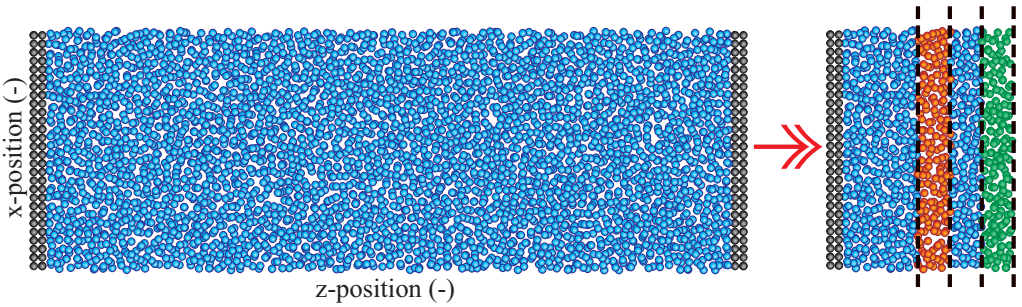

Figure 3.14: The model layout used in the coupled MD/continuum simulation. The left graph shows a nanochannel that is completely modeled with MD. The right graph shows how the same nanochannel can be modeled by using MD near the walls, while the rest of the solution is obtained with a continuum method.

tinuum domain, which shares an overlap region with the MD domain, solves the continuum equation. Note that for the coupled simulation only half of the channel has to be simulated, because unlike the pure MD simulation, in the continuum it is very easy to apply a symmetry boundary condition. The Poiseuille velocity profile is at its maximum there, so the boundary condition is simply described by $\partial u/\partial z = 0$. The other boundary condition for the continuum domain must be obtained from the MD domain.

The full MD simulation is initialised into $14 \times 8 \times 42$ lattice sites, where the top and bottom lattice sites are marked as solid atoms, composing the walls of the nanochannel. The number density of atoms in the MD domain is set to $\rho = 0.8$, which means the MD domain size is $23.940 \times 13.680 \times 71.819$ units and that the centres of the walls are separated by 70.109 units. This corresponds to approximately 23.3 nanometres in physical units (for argon). The temperature is kept constant by a Nosé-Hoover thermostat at $T = 1.2$. For the coupled simulation only the region near one of the MD walls, where the non-continuum effects are present, needs to be simulated with MD. However, the overlap region used for the Schwarz alternating method need to solve a continuum solution. Therefore, the minimum size of the MD domain is limited. In Section 2.4 it was shown that the non-continuum effects at least extend to a distance of about 8.5 units away from each wall, making this the minimum domain size. Therefore, the MD domain in the coupled simulation is initialised in $14 \times 8 \times 11$ lattice sites, where now only the bottom lattice site is marked representing a solid (atomistic) wall. This MD domain size is $23.940 \times 13.680 \times 18.810$, which effectively reduces the size and computational effort by a factor 4, while simulating exactly the same solution. As before, the cell list structure used for the efficient computation of the interatomic forces is also used for the coupling. As shown in Figure 3.14, the two top cells are used for the forcing-bin, while the center of the boundary-bin, which also consists of two layers of cells, is placed four cell distances (6.27 units) away from the centre of the forcing-bin. For the coupled simulation, the continuum domain is updated every 375 time units, which is in this case is equal to 75×10^3 MD time steps. This number was found to be a good compromise between convergence of the statistics and the computational effort. During the simulation the convergence of the solution is checked. Once it is detected that the solution has converged sufficiently, the coupled simulation is continued for an additional of 1×10^6 MD time steps. This is done in order to obtain sufficiently converged statistics to create the average profiles. The full MD simulation of the channel is run for a fixed total of 2.4×10^6 time steps. Similar to the coupled simulation, only the last 1×10^6 time steps are used for making the average profiles.

Next, the boundary conditions are discussed. As shown in the previous section, Poiseuille flow is simulated by applying a body force and solves for stationary Stokes flow with an external force for the MD and continuum domains, respectively. The MD domain has periodic boundary conditions in the x-direction and y-direction. For the full MD simulation the boundary conditions in z-direction are determined by the top and bottom MD walls. For the coupled MD simulation the MD wall is only used for the bottom side of the domain, while the continuum solution is used for the top side of the domain. The applied body force

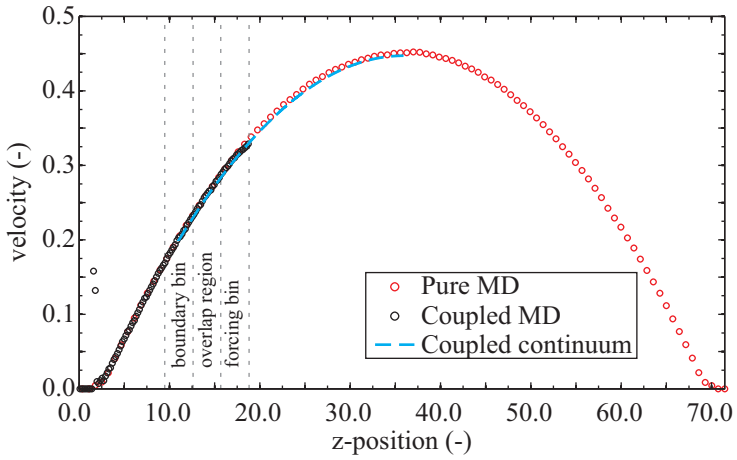

Figure 3.15: The resulting velocity profiles from the full MD simulation and the coupled simulation.

is in the x-direction and has a value of $f_{x,ext} = 0.0020$. Because the MD domain is periodic in the x-direction, the continuum domain solves a 1D equation. Therefore, only two boundary conditions are needed. The first boundary condition is taken from the average MD solution, obtained from the boundary-bin, while the other one is the symmetry boundary condition discussed above. Finally, the value for the viscosity for the continuum domain needs to be specified. Just as before, a value of $\mu = 2.0$ is chosen. However, because also a full MD simulation is performed, this value is again verified against these results.

Figure 3.15 shows the velocity profiles and density profiles obtained from the pure MD simulation and from the coupled simulation. It is evident that the coupled simulation successfully simulates the non-continuum details near the wall, while the remaining part of the profile obtained with the continuum equation are identical to those obtained with the full MD simulation. However, by using the coupled simulation, considerably less computation time was needed to obtain the same final solution, e.g. good part of the day versus only several hours. Furthermore, from the full MD simulation the value for the viscosity is extracted and yields $\mu_{measured} = 1.98$, which is almost identical to the value used for the viscosity in the continuum domain.

In this section, a coupled simulation was shown where Poiseuille flow in a nano channel is simulated. The strength of this coupled simulation is that the non-continuum effects near the wall are simulated and no expensive MD computation time is wasted on the continuum part, while the total solution is the same as the solution obtained with a full MD simulation. Although the simulation shown above is somewhat idealistic, the same principle can also be used for cases where the wall-fluid interaction is very important and can not be correctly simulated by standard continuum techniques, like phenomena such as wetting and velocity slip in microchannels [103].

3.3.3 Coupled Simulations of Flow round nano-sized particles

In the previous section it was shown how the coupled simulation can be seen as a new boundary condition for the continuum. From the continuum perspective nothing has changed; the boundary condition is still a Dirichlet boundary condition, where the value is now more accurately supplied by the communication of the MD and continuum domains. In this section the opposite is shown, namely how coupling MD and continuum enables boundary conditions for MD systems that are difficult or impossible to implement in a pure MD case.

The most used boundary conditions for MD simulation tend to be periodic boundary conditions. This is partially because these boundary conditions are very easy to implement in an MD code, and partially because this boundary condition is used to try to simulate the bulk instead of a localised problem. On the other hand, periodic boundary conditions are simply

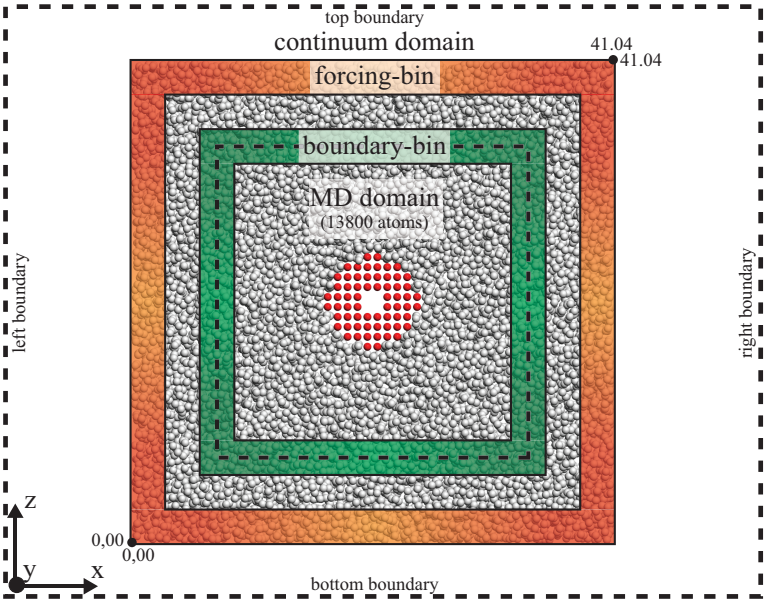

Figure 3.16: The principle of two-dimensional coupling. The MD domain is surrounded by a continuum domain. The three dimensional MD domain is periodic in y-direction and is coupled to the continuum domain using the boundary-bin and the forcing-bin in the other directions. Inside the continuum domain a two-dimensional differential equation is solved.

used because there is no good alternative. However, there are situations where non-periodic boundary conditions are essential to accurately solve the problem, e.g. the influence of *one* obstacle instead of the influence of the one and all its many periodic images. In this situation the coupling between MD and the continuum can be used to accomplish this, as is shown below.

So far only one-dimensional coupled simulations were shown, where the other two dimensions were still periodic. The next step is to remove another periodic boundary condition and to substitute it with a coupled boundary condition. This means that coupling now is done for full two-dimensional cases, where only one boundary is periodic. This principle is similar to pure continuum two-dimensional flow, where basically the third dimension is eliminated from the equations. However, for the coupling, this added dimension implies several important things. First of all, the MD domain now is surrounded by a continuum domain both in the x and z directions. Furthermore, atoms are now free to leave and enter the domain anywhere in any direction. This also means that the additional applied force inside the forcing-bin and the gradient as discussed in Section 3.2.2.1 must be applied in two dimensions. However, from another perspective, all the external boundary conditions are now continuum boundary conditions that are communicated with the MD domain. To demonstrate the 2D coupling, continuum flow around a nano-sized obstacle is simulated. Two distinctive features of the coupled simulation compared to a full MD simulation are that without the need for periodic boundary conditions the flow round a *single* nano-sized obstacle is simulated, and the external boundary conditions can be anything that is possible for the continuum. Figure 3.16 demonstrates this 2D coupling in more detail.

Up to this point it was not yet specified how the nano-sized obstacle looks like. However, because the obstacle is modelled in MD it can be any molecular structure of interest, which can be a nanotube, nanowire, or even a complex molecule like DNA. Furthermore, in the case of a geometric obstacle (e.g., a spherical particle or membrane), any "real" lattice structure of a solid material can be modelled where the material properties can be matched by adjusting the two Lennard-Jones parameters ϵ and σ. Clearly, the coupled simulation has another added benefit compared to a pure continuum simulation where the obstacle must be modelled with a no-slip (or similar) boundary condition. However, in order to show

the main principle of the two-dimensional coupling, in the following simulation a nanowire is simulated. This nanowire is modelled similar to how the walls of the nano channel in the previous Section are modelled, i.e. with atoms fixed inside a lattice. The coupled MD/continuum simulation positions this nanowire in the centre of the whole domain and in the continuum domain a uniform flow is applied, which will be enforced on the nanowire through the coupling. Figure 3.16 shows this situation, demonstrating the different bins and boundary involved. The uniform flow inside the continuum domain is generated by applying a constant velocity boundary condition, U_u, at the left, top, and bottom boundary. The right boundary boundary condition is set to an *outflow* boundary condition. This basically means that the directions of the velocity vectors are perpendicular on the right boundary and is applied by a Neumann boundary condition. The continuum solution is enforced to the MD domain inside the forcing-bin, while the new boundary conditions for the continuum domain are supplied by the boundary-bin. The centre-line of the boundary bin (the dotted line in the figure) represents the inside continuum boundary.

The MD domain is initialised in $28 \times 5 \times 28$ lattice sites, where the size of the domain is $41.04 \times 10.26 \times 41.04$ units. The nanowire has a diameter of 8.0 units which about 2.8 nanometre when using the parameters for argon. The number of atoms representing the nanowire are $N_{obst} = 384$. The number density of the fluid is taken $\rho = 0.8$, which means a total of $N_f = 13416$ atoms are representing the fluid around the nanowire. The temperature is controlled with the Nosé-Hoover thermostat to $T = 1.2$. The thermostat is only applied inside the forcing-bin. For the situation sketched in Figure 3.16 this means that the two outer layers of cells are used for the forcing-bin, while the center of the boundary-bin, which also consists of two layers of cells, is placed four cell distances (5.86 units) away from the centre of the forcing-bin. For the coupled simulation, the continuum domain is updated every 500 time units, which is in this case is equal to 1.25×10^5 MD time steps. This number was found to be a good compromise between convergence of the statistics and the computational effort. During the simulation the convergence of the solution is checked by fitting the sum of the values of the boundary conditions with an exponential function (see also Section 3.3.1). Once it is detected that the solution has converged sufficiently ($< 4 \times 10^{-3}$), the coupled simulation is continued for an additional of 5×10^5 MD time steps. As before, this is done in order to obtain sufficiently converged statistics to create the figures under constant conditions. The boundary conditions for the MD domain in the x and z-directions are specified by the continuum solution, while in the y-direction periodic boundary conditions are specified. The size of the continuum domain is set to three times the MD domain size in x-direction and to times the MD domain size in z-direction, i.e 123.12×82.08 units. Finally, the value for the viscosity for the continuum domain needs to be specified. Just as before, a value of $\mu = 2.0$ is chosen. The uniform flow velocity that is simulated is $U_u = 0.25$ units.

Figure 3.17 shows the resulting velocity magnitude contours obtained by averaging the last 5×10^5 MD time steps of the coupled simulation. The total number of time steps needed to obtain this solution was 7.5×10^6 time steps, which is equal to 60 Schwarz iterations. At that point the solution was converged sufficiently enough (relative error$< 4 \times 10^{-3}$), which basically means that the change of the coupled solution did not change more than what can be expected from the statistical noise from the MD model. The velocity contours show how the uniform flow, specified at the left, top, and bottom, flows around the single nanowire. The centre-lines of the forcing-bin and boundary-bin are displayed with the dotted lines inside the figure. The maximum velocity is reached between the nanowire and the top and bottom boundary. The velocity near the nanowire is low to very low, similar to the no-slip boundary condition. The solution is not completely symmetric with respect to cross-section (BB), which has two reasons. The first reason is the fact that the Reynolds number [20] in this case is $Re = \rho U_u D/\mu = 0.8$, while Stokes flow, which would be symmetric, assumes $Re \ll 1$. However, the MD domain does not have this restriction and any non-linear effects inside the MD domain will result in non-symmetric boundary conditions for the continuum domain. The second reason is the numerical and statistical noise in the MD domain involved in obtaining the boundary conditions for the continuum domain. One way to avoid this is to make the extracted boundary conditions for the continuum for cross-section (BB) symmetric before each Schwarz iteration. This for example, is done by Werder et al. [206]. However, doing so will also limit the type of flow that can be simulated in this coupled situation. In general, it is hard to distinguish the border between the MD solution and the continuum solution from the contours alone, meaning that the coupling method is working. However,

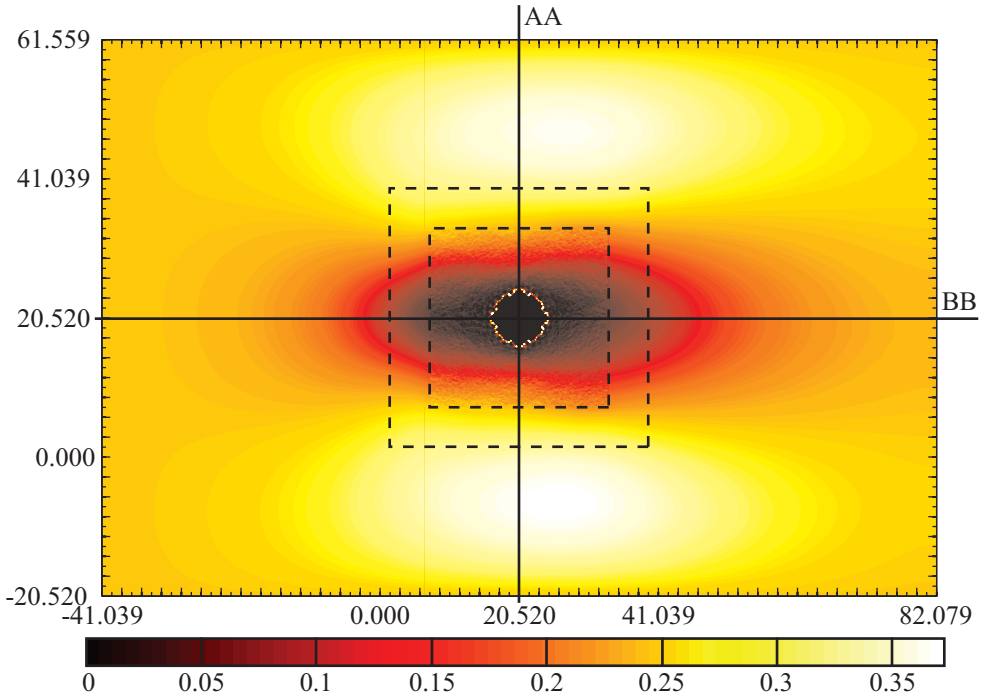

Figure 3.17: The velocity magnitude contours demonstrating the results from the coupled simulation of uniform flow round the nanowire. The two cross-section AA and BB (see text) are also shown.

in order to investigate the velocity near the nanowire more closely and to compare the values with a pure continuum solution, two cross-sections of the velocity contour plot are investigated. The first cross-section (AA) is taken from the top to the bottom boundary, positioned at the centre of the nanowire. The other cross-section (BB) is taken from the left to the right boundary, again through the centre of the nanowire.

Figure 3.18 shows the results of the values of the velocity at the two cross-sections AA (left) and BB (right). Both graphs contain three data sets. The first two data sets, the circles and the solid line, represent the coupled solution. The third data set represents the solution obtained from a pure continuum technique. This is accomplished by also simulating the nanowire (as a smooth cylinder) and the flow near the nanowire using a continuum technique. The boundary condition on the nanowire is the no-slip boundary condition, while the other boundary conditions remain the same. The velocity profiles shown in Figure 3.18 show that there is a difference in solution between the coupled situation and the pure continuum situation. This is especially noticeable near the nanowire, where in the coupled situation the value of the velocity is much lower than the one computed in the pure continuum situation. On the other hand, by definition, the value of the velocity at the left, right, top, and bottom boundary are the same in all situations. This, together with the fact that near the nanowire the velocity is lower than in the pure continuum case, this means that the value of the maximum velocity inside the whole domain is higher in the coupled situation. A possible reason for why the velocity is lower near the nanowire in the coupled situation, is discussed next.

By simulating the interatomic interactions between the fluid and the nanowire, non-continuum effects, like density variations, are simulated. Very similar to what was the case inside the nano channel simulated in Section 2.4.2, near the nanowire density variations occur up to a distance of about 8.5 units. These density variations can also mean that the local viscosity is affected. In a situation where the local viscosity near the nanowire is much higher than the value of viscosity in bulk conditions, this would mean the velocity would be smaller than expected from the case where the viscosity is everywhere the same. As is evident from figure

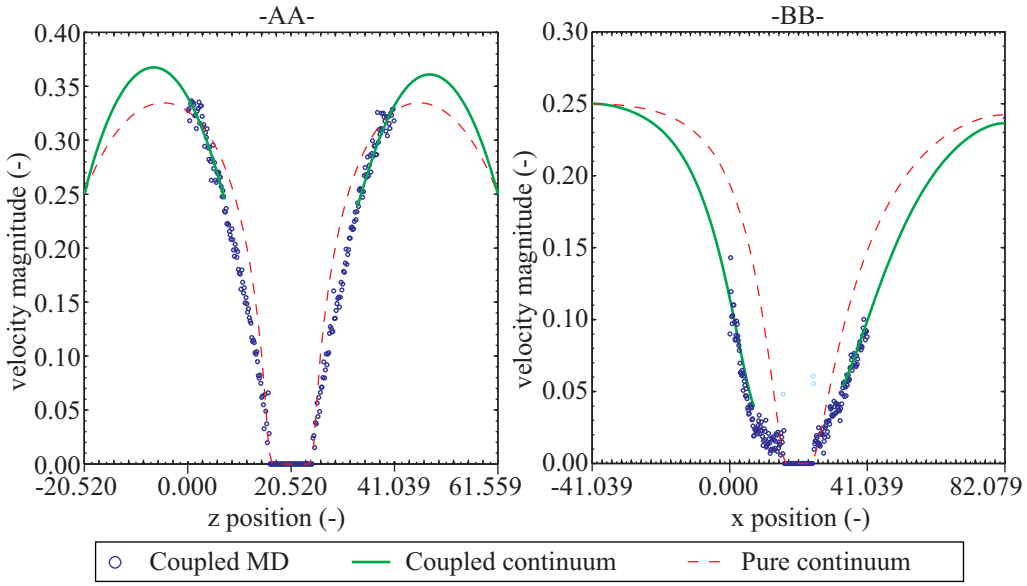

Figure 3.18: The value of the velocity taken from two cross-sections positioned at the centre of the nanowire. The coupled values are compared to a pure continuum case.

3.18 this is exactly the case, i.e. near the nanowire the velocity in the MD/Continuum coupled case shows a much lower value than in the pure continuum case. However, in order to further prove the possibility that the local viscosity is the cause of the lower velocity, another type of coupled simulation is performed next. In this coupled simulation, the nanowire is replaced with a nano structure as shown in Figure 3.19.

The nano structure has the purpose to make a small duct through which the fluid should flow. The height of the duct is 10 units, which is about 3.4 nanometres when using the parameters for argon. This means that the density variations are present inside the whole duct. As before, the continuum boundary conditions simulate uniform flow. If the viscosity near the MD structure is not affected by the density variations, the velocity inside this duct (shown left in figure 3.19) should more or less be equal to the pure continuum situation (shown right in figure 3.19). The structure in the pure continuum situation is modelled with a smooth object, i.e. the boundary condition at the nano structure is set to the no-slip boundary condition. However, in order to test the effect of the value of the viscosity in the duct, this area can be modelled using a different value of the viscosity, μ_{duct}. The two continuum cases that are simulated use a viscosity of $\mu_{bulk} = \mu_{duct}$ =2.0 units and $\mu_{duct} = 12.0$ units (i.e. $6 \times \mu_{bulk}$). From the simulations, the velocity profile along the cross-section (AA) is obtained, which should show a difference between the two cases simulated. The other simulation conditions are exactly the same as before.

Figure 3.20 shows the velocity profiles obtained from the coupled simulation (top right), the two pure continuum simulations (top left), and the the continuum velocity contours inside the domain when the viscosity inside the duct is taken six times higher than the bulk (bottom) . First the pure continuum solutions (top left graph) are discussed. As is evident, a higher viscosity inside the duct results in a lower velocity inside the duct. The main reason for this is the fact that a higher viscosity in the duct means it is less favorable for the fluid to go through the duct. The velocity contours in this situation (the bottom graph) show how almost no fluid is going through the duct, but instead flows around the nano structure. The velocity profile from the coupled MD/continuum solution shows that, in the coupled situation, also almost no flow is present inside the duct, similar to the higher viscosity case for the pure continuum simulation. Note that the value of the velocity inside the MD duct is within the statistical noise of the MD simulation. Hence, any detailed observation of the coupled solution must be done sceptically. However, the coupled results compared to the pure continuum results clearly show that when the value of the viscosity inside the duct

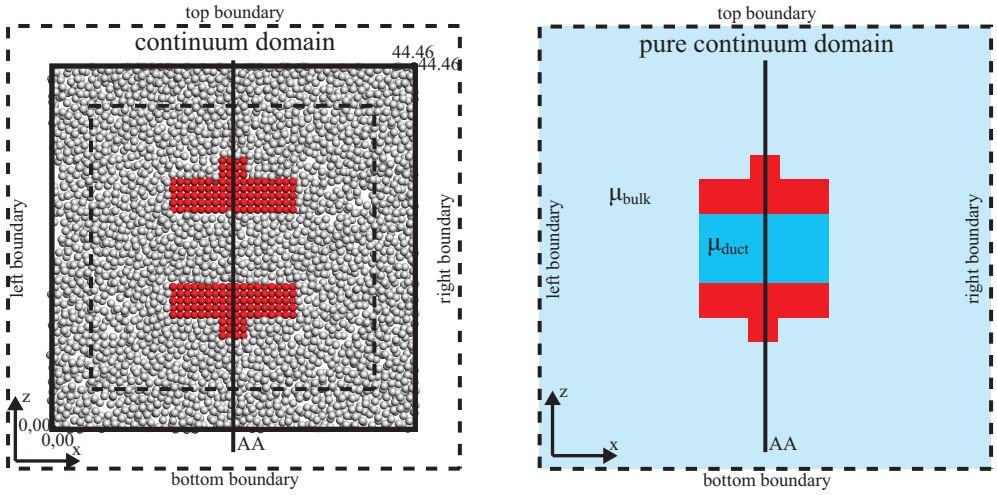

Figure 3.19: The model of the nano structure, which effectively act as a duct, that is used to test the influence of the density variations on the viscosity. The left figure shows the coupled MD/Continuum model, while the right figure shows the pure continuum model where the (continuum) viscosity is set differently than the bulk value inside the duct (μ_{bulk} versus μ_{duct}).

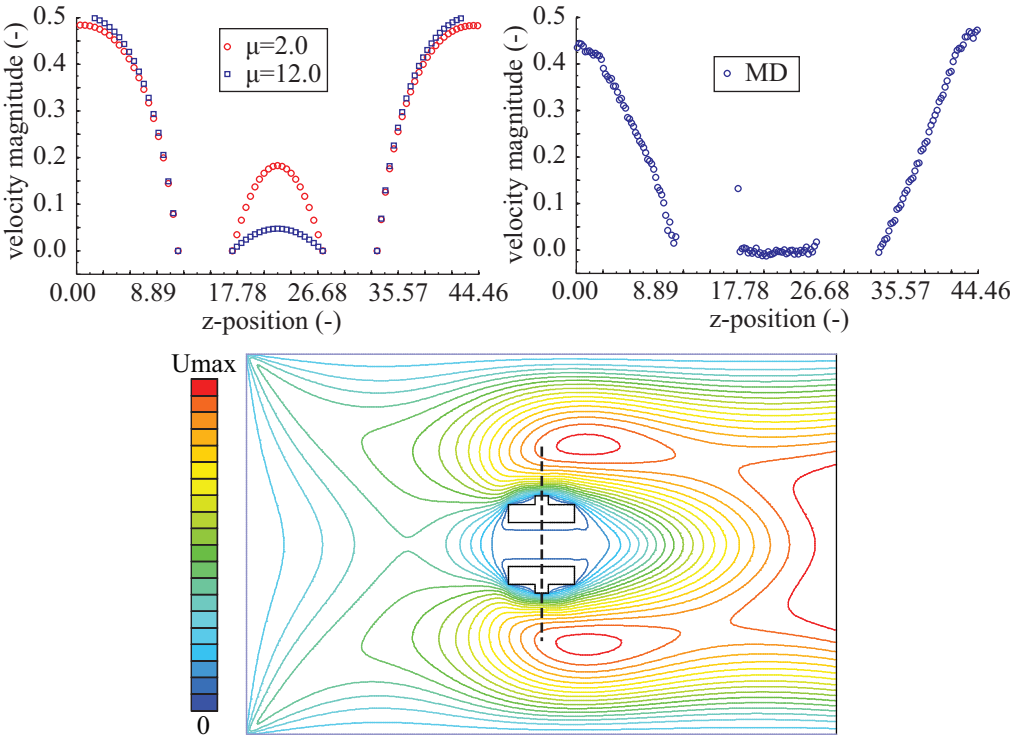

Figure 3.20: The velocity profiles obtained from the pure continuum situation (top left) and from the coupled MD/Continuum model (top right) when simulating the nano structure that creates a small duct. The bottom graph shows the continuum velocity contours inside the domain when the viscosity inside the duct is taken six times higher than the bulk.

is much higher than the bulk value, both simulations solve the similar solutions. In other words, the influence of the density variations on the value of the velocity near nano-sized obstacles can be explained by the fact that the viscosity is much higher near these nano-sized obstacles.

This concludes this section where the flow around a single nanowire is investigated using an two-dimensional MD/Continuum coupled simulation. The simulation showed that the density variations near the nanowire influences the velocity near the nanowire and therefore the velocity contours inside the whole domain. In a separate coupled simulation it was shown how the local viscosity near the MD obstacles can become very high and therefore influences the flow. Therefore, the strength of this coupled simulation is the fact that the (continuum) flow around a single obstacle is simulated, where the effects of the molecular structure of the obstacle can be simulated very accurately. From another perspective, the coupled simulation also enables the specification of any continuum boundary condition for an MD domain.

3.3.4 Coupled Simulations Using Water / Boundary Force for Water

The coupled simulations so far, simulated a liquid that consisted of single non-polar atoms, i.e. argon. It was shown that this type of coupled simulation can be very helpful to investigate flow problems where the most important effects, e.g. wall effects, are happening at the scale that is easily accessible to an MD simulation, while there is a macroscopic effect that would not be easy to be simulated with an MD simulation alone (due to the computational time needed). In other words, these simulations with single non-polar atoms can give good insight of the behaviour of liquids at very small scales (i.e. nanometre scale). However, especially for biological applications, another type of atoms/molecule plays a very important role, namely water. As shortly explained in Section 3.2.3, the methods used for the coupling of water are very similar to those used for the coupling of argon. However, before a successful coupling can be achieved, some problems first need to be solved.

One of the problems to be solved is the elimination of the unnatural variations in the MD results near the MD/continuum boundary as a result of the incomplete interatomic interactions. As explained before, in the case of argon, these incomplete interactions could be artificially supplied by adding an additional boundary force to all the atoms that are near the MD/continuum boundary. This boundary force is a function of the distance between the atom and the boundary and can be found by a separate simulation or by iteration (see Section 3.2.2.3 and Kotsalis et al. [110]). However, for water the situation is trickier, because now a molecule, rather than a single atom, must be considered. One attempt to perceive a boundary force for the use of molecular dynamics simulations of water, was done recently by Kotsalis et al. [109]. Here the density/boundary force iteration method the same author developed before [110], is applied to water. The iteration is performed between the obtained and desired density profile to find a single boundary force that can be applied to the centre of mass of each water molecule. Using this technique, the density variations in the density profile of water can completely be eliminated. However, it is also shown [109] how the probability distribution of the orientation of the water molecules near the MD/continuum boundary is very much different than the distribution somewhere else in the domain. In other words, if nothing is done to prevent so, the water molecules have a certain preferred orientation near the boundary. This preferred orientation can be the result of other missing interactions (e.g. long-range interactions), the influence of the specular wall used as the MD domain boundary, or simply because the boundary force is applied to the centre of mass. To investigate the influence of the position where the boundary force is applied, an adaptation of the density/boundary force iteration method is developed next. The difference between this method and the method described in [109] is the fact that now for each atom inside the water molecule, a separate boundary force is obtained and applied at the position of the atom itself (instead of the centre of mass of the molecule). Of special interest is the question whether this has any effect on the preferred orientation of the water molecule near the MD/continuum boundary.

To find the boundary force function, the model layout shown in Figure 3.21 is used. The MD domain size is $3.94 \times 3.94 \times 5.71$ nanometres and the total number of atoms simulated is $N = 8480$. The bottom boundary condition of the MD domain is given by an atomistic

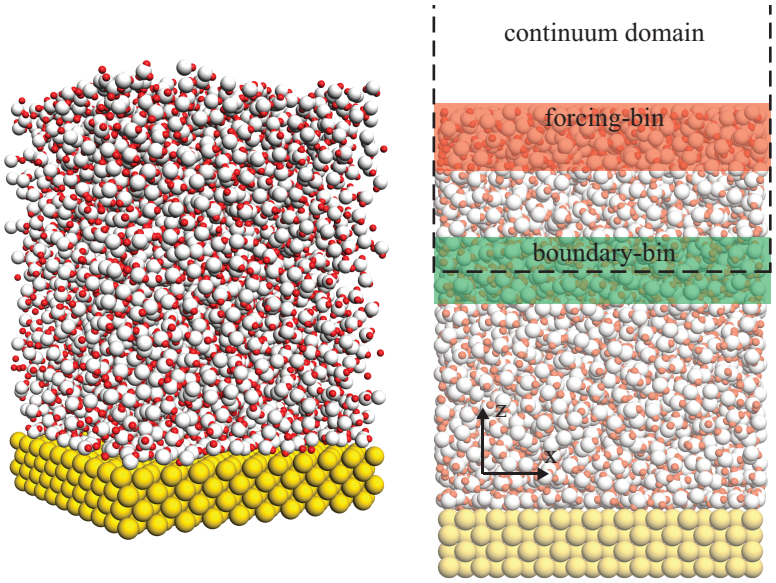

Figure 3.21: The model used in the coupled simulation inside a nano channel filled with water. Only the atoms/molecules near the wall need to be considered with MD, while the rest of the channel can be simulated with continuum methods. The wall consists of 800 atoms positioned inside a FCC lattice. The liquid is represented by 2560 water molecules, making the total number of atoms simulated, $N = 8480$. The MD domain size is $3.94 \times 3.94 \times 5.71$ nanometres.

wall consisting of 800 atoms with a fixed position in a FCC lattice. The liquid (water) is represented by 2560 water molecules, which means, for the volume available, the initial density is $1000\,kg/m^3$. The water model used for the interactions is the TIP4P/2005 water model [2], for which the parameters can be found in Section 2.3.6. The top boundary condition of the MD domain is a specular wall, while the four remaining boundary conditions of the MD domain are chosen periodic. The continuum domain is coupled with the top boundary of the MD domain. In order to accomplish the coupling, two bins are specified. The first bin, i.e. the boundary-bin, is the bin inside the MD domain where the macroscopic variables for the continuum boundary conditions are extracted from the MD simulation results. The other bin, i.e. the forcing-bin, is the bin where the continuum solution is enforced onto the MD domain. The remaining boundary conditions for the continuum domain are set to zero everywhere, i.e. no flow is present or enforced. The area near the MD domain boundary (of which the forcing-bin is part of), is the area where the boundary force function is applied to the atoms inside the MD domain. The cutoff radius used for the Lennard-Jones interaction potential in the following simulations is $r_c = 0.9$ nanometres. On the other hand, the Coulombic interactions are evaluated using a PPPM technique (see Section 2.3.4) and therefore are not truncated. However, because the MD domain is not periodic in three dimensions, something must be done to prevent wrong interactions in the z-direction. The technique used here, is the ELC-method [11, 99], where besides an artificial vacuum on top of the MD domain, also an additional correction, the so-called electrostatic layer correction (ELC) is applied. Please note that this not mean that the missing (long-range) interactions from the continuum domain are taken care of, because the two-dimensional PPPM techniques only limits the unwanted periodic replications from the interatomic interactions. Therefore, the boundary force function still needs to provide these missing interactions. In order to evaluate the (total) boundary force function efficiently, this function is only evaluated up to a distance of the cutoff radius away from the MD boundary. Nevertheless, this does mean that the total boundary force function includes both Lennard-Jones interactions and Coulombic interactions.

Before a simulation is started, the MD system is first initialised by "melting" the system

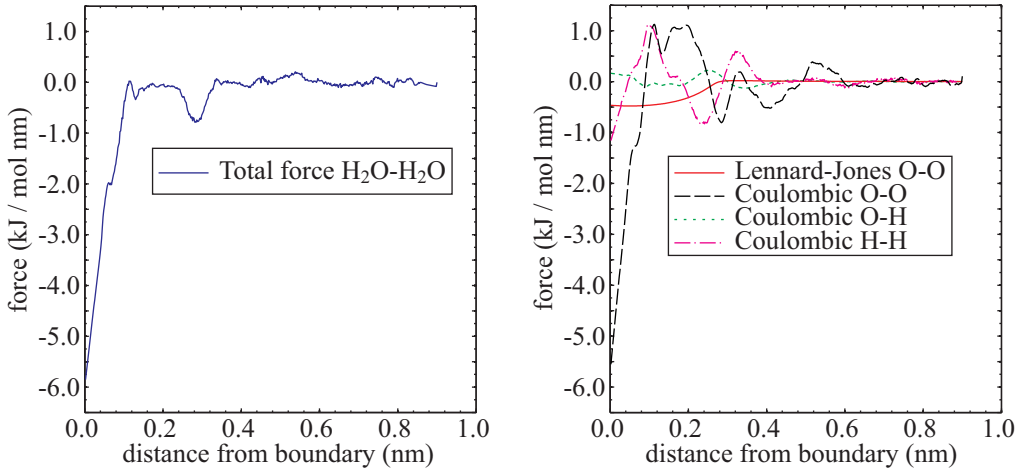

Figure 3.22: The boundary force function obtained from the iteration simulation. The left graph shows the function when the force is only applied to the centre of mass of the water molecule, while the right graph shows the boundary force applied to the individual atoms inside the water molecule.

by applying a Nosé-Hoover thermostat with a target temperature $T = 293\,K$. After the initialisation, the simulation to find the boundary force function is run for a total of $3.2\,ns$ with a time step of $2.0\,ps$, i.e. 1.6×10^6 time steps. As mentioned before, two different types of boundary force functions are investigated. The first one is where the force is only applied to the centre of mass of each water molecule. The second one is where a force is applied to each atom inside the water molecule. In both cases, the continuum domain boundary conditions are set to zero everywhere, i.e. no flow is present or enforced. In each simulation, the density profile near the MD/continuum boundary is collected in 500 bins. These 500 bins are positioned between the boundary and a cutoff radius away from the boundary. Data is collected for 4×10^4 time steps after it is compared with the target density. The difference between the obtained density profile and the target value determines the difference in the boundary force function [110]. The iterations continue by collecting the density again, after changing the boundary force function accordingly. In the case when only one boundary force function at the centre of the water molecule is simulated, the update of the boundary force function is straightforward, i.e. all the differences are applied to the one function. However, in the case where multiple boundary force functions are used, the update of the boundary force function also must be distributed amongst the multiple functions. This also means that there is no unique solution and consequently, several combinations of different boundary force functions applied to the individual atoms can result in the same target density profile. Therefore, in order to limit the complete randomness of the resulting boundary force functions, first a separate simulation is performed. This simulation is similar to what was done in Section 3.2.2.3, where the missing Lennard-Jones interactions for argon are determined from a complete periodic MD simulation. However, for water this means that four different interactions need to be determined. These are, the Lennard-Jones interaction between the O atoms and the Coulombic interactions between each of the O atoms, the O and H atoms, and each of the H atoms. Once these boundary functions are determined, they can be optimised using the iteration method, where the updating of the boundary force function can be distributed amongst the four different types of interaction. However, it was soon discovered that the Coulombic interaction between the O atoms was the most important force. Therefore, the iteration is only applied to this force.

Figure 3.22 shows the results of the simulations after the iterations. The left graph in Figure 3.22 shows the obtained boundary force function when this force is only applied to the centre of mass of the water molecule, while the right graph in Figure 3.22 shows the obtained boundary force functions that are applied to the individual atoms inside the water molecule. When comparing the left and right graph, two observations are most obvious.

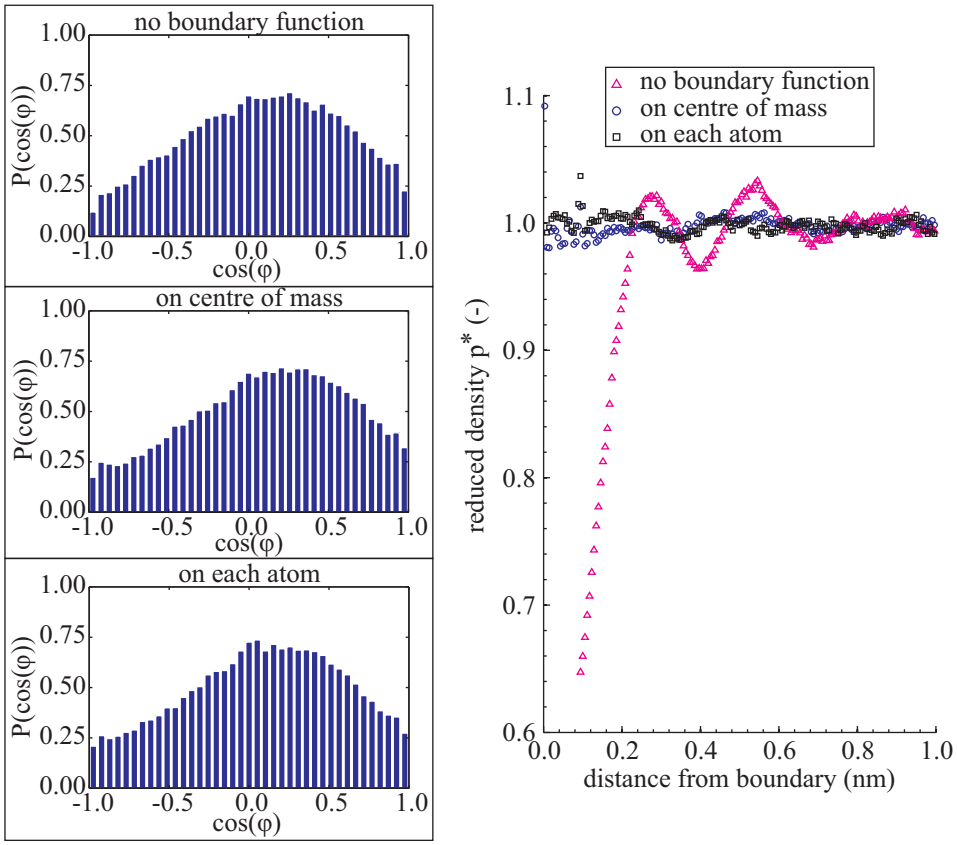

Figure 3.23: The probability distribution of the orientation of the water molecules near the MD/continuum boundary and the resulting normalised density profile ($\rho^* = \rho/\rho_{bulk}$). The left graphs show the probability distribution for the case when no force is applied, when the force is applied to the centre of mass of the water molecule, and when the force is applied to each individual atom inside the water molecule. The right graph shows the density profiles near the MD/continuum boundary for all cases.

The first being the fact that, as mentioned before, the Coulombic interaction between the O atoms is the most predominate force present. The second observation is that both boundary force functions show a large negative value very near the boundary (< 0.2 nm). This implies that the boundary force function pulls atoms strongly towards the MD/continuum boundary in order to obtain the target density. Next the boundary force function obtained for the centre of mass of the water molecule (left graph) is compared to this same function obtained by Kotsalis et al. [109]. The comparison reveals a large difference between the two obtained boundary force functions. One possible explanation for the big difference, could be the water model that is used in the simulations. Kotsalis et al. [109] used the SPC/E water model (see Section 2.3.6 and [16]), while here the TIP4P/2005 water model [2] is used. In order to verify this statement, a separate simulation using the SPC/E water model was performed. However, the results (not shown here) still revealed a large difference between the two boundary force functions. This also strengthens the argument that, unlike in the case of argon, no unique (single) boundary force function can be found for water as such. However, by definition of the iteration method, in all cases the density profile is equal to the target density. This, and the results of the orientation of the water molecule near the boundary, are shown next.

Figure 3.23 shows the results of probability distribution of the orientation of the water molecules near the MD/continuum boundary and the resulting normalised density profiles ($\rho^* = \rho/\rho_{bulk}$). The left graphs show the probability distributions for the case when no

boundary force is applied, when the force is applied to the centre of mass of the water molecule, and when the force is applied to each individual atom inside the water molecule. The right graph shows the density profiles near the MD/continuum boundary for all the cases. From the results of the orientation of the water molecule near the MD/continuum boundary it can be concluded that in all three cases the water molecule has a preferred orientation. Also, by comparing the two cases when a boundary force is applied, the position where this boundary force is applied, does not have a significant effect on the preferred orientation of the water molecule. Further, by comparing the case where no boundary force is applied and the cases where it is applied, it must be concluded that the boundary force itself does not have a significant effect on the orientation of water molecule at all. In other words, the applied boundary force is not the (main) reason why the water molecule has a preferred orientation near the MD/continuum boundary. Other effects, like long-range interaction and the specular wall, are more likely to cause the preferred orientation. Note that the probability distribution shown in [109] is very similar to the results obtained above. However, in all cases (not shown here) the probability distribution of the orientation of the water molecule recovers the correct uniform distribution, i.e. no preferred orientation, as the distance from the boundary is increased. Therefore, the (incorrect) preferred orientation of the water molecule is only a local effect. Next the obtained density profiles are discussed. The density profile in the case when no boundary force function is applied, shows a steep decrease in density very near the boundary (> 0.2 nm), while the remainder of the profile only shows several gentle variations of the density before bulk density conditions are obtained. The density profiles in the case when a boundary force is applied, either on the centre of mass or on each individual atom, shows the bulk density everywhere. This means the density/boundary force iteration is doing what it is designed for, i.e. limiting density variations.

This concludes the coupled simulations where it is shown how the boundary force function can be used to eliminate the unnatural density variations near the MD/continuum boundary. It is also shown that the water molecules do not have the right probability distribution of the orientation near this boundary. However, in the remainder of this section it will be shown how a coupled simulation of water can successfully be used for Poiseuille flow inside a nano-sized channel.

The procedure for the coupled simulation of Poiseuille flow of water inside a nano-sized channel is very much similar to the coupled simulation of argon discussed in Section 3.3.2. For this purpose, the same model shown in Figure 3.21 can be used. This means that effectively only the flow near the atomistic wall is considered with MD, while the rest of the flow is described by a continuum equation. Again, note that for the coupled simulation, only half of the channel has to be simulated, because unlike the pure MD simulation, in the continuum it is very easy to apply a symmetry boundary condition. This means for the continuum that the boundary condition is simply described by $\partial u_x/\partial z = 0$. The other boundary condition for the continuum domain must be obtained from the MD domain from the boundary-bin. The z-coordinate of the middle of the nano-sized channel, z_2, is chosen to be $1.5\times$ the total MD domain height, i.e. $z_2 = 8.57\,nm$. The z-coordinate of the centreline of the boundary-bin, i.e. the other continuum boundary, is positioned at $z_1 = 4.28\,nm$. The z-coordinate of the centreline of the forcing bin is $z_{MD} = 5.43\,nm$. In the MD domain, the Poiseuille flow is generated by a body force acting on each water molecule, while the continuum solves for stationary Stokes flow with an external force. The value of the external force is determined from the fact that the velocity inside the nano-sized channel should not be too high, however large enough in order to collect meaningful statistics in a relatively short time. A value of $f_b = 5.98 \times 10^{-4}\,J/\,(mol \cdot nm)$, which would mean the velocity in the middle of the channel is about $50\,m/s$. The temperature is kept constant by a Nosé-Hoover thermostat at $T = 293\,K$. The last parameter that needs to be specified for the flow, is the value of the viscosity. As mentioned before, the value of the viscosity is a result of an MD simulation rather than a specified parameter, which is the case for a continuum computation of flow. For the successful coupling of an MD domain and a continuum domain it is therefore important to match the value of the viscosity. As shown in Section 2.4.5, several commonly used water models, like SPC/E [16] and TIP4P [102], underpredict the value of the viscosity by as much as 30-50%. It was also shown that the water model TIP4P/2005 [2] does give satisfactory value of the viscosity, comparable to experimental values of the value of viscosity for water. Therefore the coupled simulation next, will use this water model. However, to show what will happen if the viscosity is not matched, a simulation with the TIP4P water model is also performed. Finally, a third simulation will be performed. This simulation is

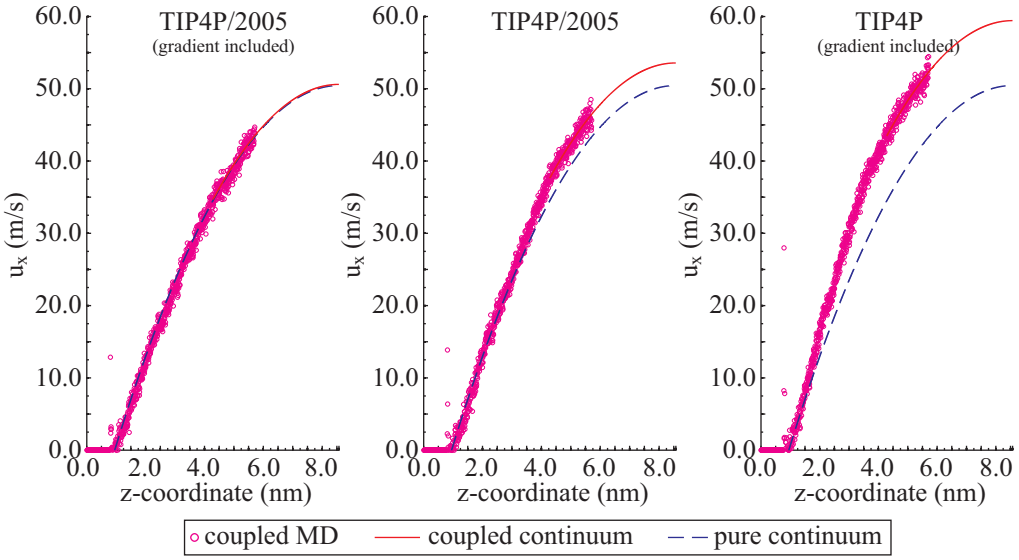

Figure 3.24: The velocity profiles (u_x) obtained from three different coupled simulation (see text). The results show a clear dependency on the used water model and the inclusion of the velocity gradient.

investigating the influence of adding the velocity gradient inside the forcing-bin (see Section 3.2.2.1 for details). If no gradient is included inside the forcing-bin, the velocity profile shows a plateau. However, this plateau will also influence the local shear and therefore also changes the surrounding velocity profile and therefore the values of the velocity used for the coupling. The simulations run for a total of 1.6 ns, with a time step of 2 ps, therefore 8×10^5 time steps are simulated. The macroscopic variables, like density and velocity, are sampled for the last 1.0 ns of the MD simulation where they are collected in 1000 datapoint bins along the z-direction (height of the channel). The boundary force function that is used here, is the function that is applied to each individual atom (see above for details). However, as also noted by Kotsalis et al. [109], the boundary force function must be adapted to compensate for the influence of the flow near the boundary. Therefore, during the simulation the iteration between the local density profile and the boundary force function is continued, in order to prevent the unnatural variations of the density near the boundary. The parameters for this iteration remain the same.

Figure 3.24 shows the velocity profiles obtained from the three coupled simulations. The left graph in Figure 3.24 shows the velocity profile obtained from the simulation where the TIP4P/2005 water model is used and the velocity gradient inside the forcing-bin is included. The middle graph in Figure 3.24 shows the velocity profile obtained using the same water model, however without the velocity gradient included. The right graph in Figure 3.24 shows the results obtained with the TIP4P water model, including the velocity gradient. In each of the graphs, the small squares represent the binned data points obtained from the MD simulation. The solid line in the graphs is the corresponding to the coupled continuum solution. This solution is obtained by solving the stationary Stokes equation with the external force where the velocity boundary condition at z_1 is extracted from the MD results. The dotted line inside the graph is a best-guess complete continuum solution of the Poiseuille flow inside the nano-sized channel. In order to obtain this curve, only the stationary Stokes equation with the external force is solved for. This means that the position of the boundary condition where the no-slip condition is valid must be guessed/assumed. Investigating the velocity profile obtained from the coupled MD/continuum solution, the z-coordinate of this boundary condition is, $z_{no-slip} \approx 1.05\,nm$. When comparing the three results, two main differences can now be observed.

The first difference is caused by the inclusion of the macroscopic velocity gradient inside the forcing-bin in the MD domain. When this gradient is included, all the solutions overlap.

However, when the gradient is not included, the coupled solution is shifted up with respect to the complete continuum solution. The reason for this is the plateau inside the forcing bin and the resulting incorrect shear. When no gradient is present inside the forcing-bin, it means (in this case) that the average velocity on the left side of the forcing-bin is too high with respect to the continuum solution. This error in value is propagating towards the surrounding bins/cells and therefore also influences the value inside the boundary-bin. Of course, the value of the velocity inside the boundary-bin is used to update the continuum solution, which now is higher than expected. This means that after the next Schwarz iteration, an even higher velocity is applied inside the forcing-bin and the discussed sequence of events further amplifies the problem. On the other hand, the values of the velocity near the wall ($z < 3.0$ nm) do not show large differences, which means that the incorrect applied shear and consequences for the coupling is only local (near the MD/continuum boundary).

The second difference is caused by the selection of the water model that is used in the MD domain. Comparing the left and right graph reveals what happens with a coupled simulation where the value of the viscosity is not matched. Inside the MD domain, the molecules are set into an average motion by applying a body force to them, which for the two simulations is the set to the same value. However, the TIP4P water model does not simulate the correct value of the viscosity and underpredicts the experimental value at $T = 293K$ by 45.2%, which is much worse that the TIP4P/2005 water model, where the value is underpredicted by 8.2%. For the coupled simulation this means that the following sequence of events takes place. 1) The velocity that is enforced inside the forcing-bin is set equal to the value obtained from the continuum equation of motion that uses the continuum (experimental) value of the viscosity. 2) The body force applied to the TIP4P water model results in a higher than expected velocity inside the boundary bin, which effectively means that the velocity inside the forcing-bin is suppressing the flow inside the MD domain. 3) The first Schwarz iteration is taking place, where the continuum solution is updated with the higher than expected velocity but with the continuum viscosity. 4) The velocity that now is enforced inside the forcing-bin is higher than expected from the pure continuum velocity curve, but is lower than expected from a pure MD simulation with the TIP4P water model. 5) The inconsistency of the value of the viscosity results in a steeper velocity curve near the wall because the value of the viscosity is lower than the continuum one. At the same time, the continuum solution is preventing the curve to become steeper near the MD/continuum boundary. This is evident in the velocity profile at $z \approx 4.0$ nm. This sequence of events results in a velocity profile that is stretched and non-parabolic compared to the complete continuum solution. This also means that the coupled simulation where the wrong value of viscosity is used/simulated, produces an incorrect solution.

As a final verification, the value of the viscosity is extracted from the three different velocity profiles. This is done using the method explained in Section 2.4.5 where the second derivative of the velocity profile with respect to the height of the channel is compared with the applied body force. The values of the viscosity are $0.999 \pm 0.08 \times 10^{-3}\, Pas$, $0.962 \pm 0.08 \times 10^{-3}\, Pas$, and $0.504 \pm 0.02 \times 10^{-3}\, Pas$ for the velocity profiles in Figure 3.24 from left to right, respectively. The experimental value of the viscosity at the simulated temperature is $\mu_{exp} = 1.01 \times 10^{-3}\, Pas$. These obtained values are similar to those obtained in Section 2.4.5 for the different water models.

3.3.5 Coupled Simulations of Thermophoretic Flow

In the previous sections successful coupled simulation are shown where the velocity field played a crucial role. However, as mentioned before, the principles of the coupling of domains can also be applied to other macroscopic variables. One of such variables that is accessible in both the MD domain and the continuum domain is the temperature. In MD the value of the temperature can be controlled by a thermostat, which generally is set constant and applied for the whole MD system (see Section 2.3.3). However, in this section a simulation is performed where only in the forcing-bin the temperature is set to a certain value while the remainder of the MD system is not thermostatted, i.e. the atoms in this area are allowed to react to the thermostatted atoms. Similarly to what was done before, the value for the temperature inside the forcing-bin must be obtained from a continuum domain. This type of simulation effectively is able to create a temperature gradient inside the two domains. In Section 2.1.3 it was shown how temperature gradient give rise to a

phenomenon called thermophoresis, i.e. the motion of micro-sized and nano-sized particles due to a temperature gradient. Furthermore, the concept of charge is very easily applied to atoms, while in the continuum this concept is easier to handle with the specification of a charge density. Nonetheless, in both domains there are ways to deal with the charges and as discussed in Section 2.1.2, very useful and interesting flows can be generated using the charges which, for example, results in an electric double layer. Therefore, in this section a coupled MD/Continuum simulation is performed that *qualitatively* studies a single nano-sized particle inside a temperature gradient field. At the end of the section it is shortly explained how the same principle can also be applied to a charged particle that is surrounded with charged ions. However, this section mostly has the purpose to illustrate some recommendations for further research, which are also discussed in Section 5.2.

First the behaviour of one single particle inside a temperature field is investigated. Previous research that used molecular dynamics to investigate this [24, 152, 158, 162, 192] concentrated on two species of different mass, interaction strength, and/or different concentration. In all these cases, as expected, one of the species moved to the hot side, while the other moved to the cold side. Although research is done using more realistic materials like, for example, liquid heptane/benzene mixtures [152], where the (pure) MD simulation obtained results were comparable to experimental values, there is still a lack of a simple physical overall explanation of thermophoresis in liquids [192] and MD simulation can definitely help in this. These (pure) MD simulations were performed with the aid of non-periodic boundary conditions, but they were not coupled to another domain. The coupled simulation, as proposed next, has the added advantage that the MD domain is surrounded by a controllable continuum domain. This opens two new options. The first one is the fact that now continuum boundary conditions can be specified and applied to the MD domain that were impossible before. The second option is that with the coupled simulation the phenomenon can be investigated at length scales that are accessible to experiments. This would be very computational expensive if the simulation was done with a pure MD code.

To illustrate the coupled simulation, the thermophoretic motion of a single nano-sized particle inside a liquid is investigated. This is accomplished by modeling the particle in a similar way as was done for the atomistic walls in Section 2.4.2 or the nanowire in Section 3.3.3. The particle is placed in the centre of the MD domain, while all sides of the MD domain are coupled to a continuum domain. This means that the set-up of the coupled simulation is identical to the situation displayed in Figure 3.16. The only difference are the type of boundary conditions that are applied and exchanged to and from the continuum domain and MD domain, which now are values of temperature rather than velocities. The particle is not allowed to move, but during the simulation all the forces exerted on the particle are collected and averaged for every direction (x, y, and z). If the particle does not experience any temperature difference, the averaged forces are zero. Therefore, any change from this is due to the effects of the temperature difference. The sign of the force indicates whether the particle wants to move to the hot or cold side. To investigate whether the size of the particle has any influence on this force, two sets of simulation with different diameters of the particle are performed. The first set simulates a particle with diameter, $D_{obst} = 8.5$ units, while the second set uses, $D_{obst} = 13.5$ units. This corresponds to a nano-sized particle of approximately 2.9 nm and 4.6 nm, respectively. The MD domain size is determined from the fact that the domain should be large enough to position the particle inside the domain and where the density variations are away from the continuum boundary. In this case this means the MD domain size is chosen as $27.36 \times 13.68 \times 27.36$ units for the small particle and $34.20 \times 13.68 \times 34.20$ units for the larger particle. As before, the value of the density is set to $\rho = 0.8$ units. The temperature difference is applied in z-direction only, induced by the continuum boundary conditions on the top and bottom boundary. The values of the temperature on the left and right continuum boundary are determined by interpolating between the top and bottom values of the temperature, i.e. the values versus the coordinate do follow a line with a constant gradient. The flow conditions for all continuum boundary conditions are set to *no-flow*, i.e. the velocity is set to zero. In all cases the simulation is performed for 2.4×10^6 time steps, while the statisitics used to make the graphs uses the last 1.6×10^6 time steps.

Figure 3.25 shows the results from the simulations. The curve in the figure shows how the total force on the particle changes when the applied temperature gradient is changed. The circles show the results for the small particle ($D_{obst} = 8.5$ units), while the squares show the results for the larger particle ($D_{obst} = 13.5$ units). The lines are linear fits to the data points.

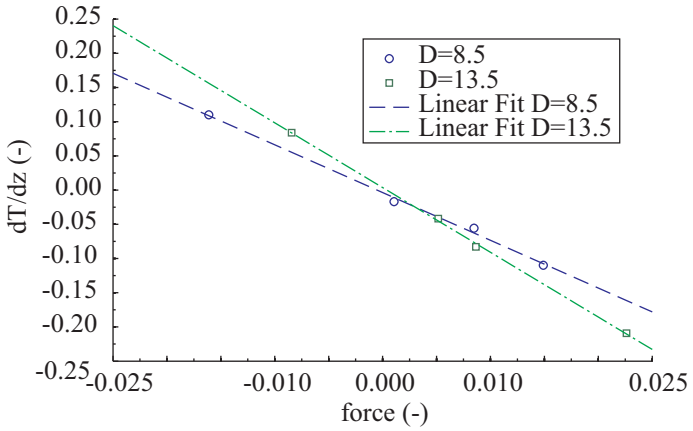

Figure 3.25: The results from the coupled thermophoretic simulation. The curve shows how the total force on the particle changes when the temperature gradient is changed. The circles show the results for the small particle ($D_{obst} = 8.5$ units), while the squares show the results for the larger particle ($D_{obst} = 13.5$ units). The lines indicate the best linear fit for the data points.

Two main conclusions can be drawn from the results. The first one is the fact that there is a small difference between the forces on the small and larger particle when they experience the same temperature gradient, where the larger particle experiences a larger force. The second conclusion is the fact that in all cases the force has an opposite sign than the temperature gradient. This indicates that in all cases the particle wants to migrate towards the cold side of the domain. This also means, that at least for the cases investigated here, there is no sign-change of the Soret-coefficient (see Section 2.1.3). This is the same conclusion derived by Bordat et al. [24], where it is explained that this seems to be case with all simulations that use the Lennard-Jones 12-6 potential applied to argon atoms. As a final verification, another simulation is performed that employs a larger cut-off radius. The cut-off radius is increased from $r_c = 2.5\sigma$ to $r_c = 4.0\sigma$ and the difference in the total force is measured. The results from this simulation showed that the total force on the particle changed (increased) by as much as 20%. This again indicates (see Section 2.4.4) that the commonly accepted cut-off radius of 2.5σ is too small for these kind of applications.

At this point, it is easy to see how the same type of simulation could also be performed with a charged particle in the centre of the domain, while this particle is surrounded with it co- and counter-ions. As shown in the pure MD case by, for example, Freund [64], large differences from the pure continuum methods can be observed at the nanometre-scale. A coupled simulation where the particle and its electric double layer is simulated with MD, while the surrounding liquid is modelled with a continuum method will solve this problem. However, as a final remark of this chapter, one flaw in the simulations so far must be pointed out. Namely, that all simulation so far concentrated on the flow round, or phenomenon near, a *static object*. In other words, the MD domain remained at the same position inside the continuum domain and did not change size during time. Therefore, to prevent the particle of interest to move outside the static domain, the particle was pinned to its position. For the realistic simulation of thermophoretic motion, electrokinetic flow, or any situation where the part of the domain that cannot be simulated with a continuum method, this part should have the freedom to move. One intermediate solution to solve this problem is to still keep the particle inside the centre of the MD domain, but at the same time move the entire MD domain (and its atoms) with the amount that the particle should have moved when it would not be pinned. However, this also means that the boundary conditions applied inside the forcing-bin must change accordingly because of the new position inside the continuum domain. Therefore, further research should point out whether these kind of simulations can still be performed efficiently using the alternating Schwarz method.

This ends this section and Chapter 3. In this chapter it was explained how the Schwarz alternating method can be used to couple an MD domain and a continuum domain. This

method, which is introduced in the first part of this chapter, works on the principle that inside an overlap region, both the MD domain and the continuum domain are solving the same result. The coupling is achieved by the correct specification of boundary conditions for each domain, obtained from the other domain. The second part of this chapter concentrated on this and it is explained how especially the specification of the boundary conditions in the MD domain require the most attention. The third and last part of this chapter showed several coupled simulation results. For example, here it was shown that the coupled simulations can successfully be used to very accurately model wall-fluid interactions with MD, while the remainder of the domain is efficiently computed with a continuum method. However, the next chapter will demonstrate a different kind of coupling between molecules and the continuum, which is very efficient to study the behaviour of polymers.

Chapter 4

Flow Injection of Polymers into Nanopores

In this chapter * the flow injection of polymers into nanopores is discussed. We use a mesoscale simulation to measure the strength of the velocity flux needed to push a polymer into a narrow channel. We find excellent agreement with the prediction by Sakaue et al. [167], based on a de Gennes blob model of the polymer, that the critical velocity flux for translocation depends linearly on the temperature, but is independent of the length of the polymer chain or the width of the channel.

4.1 Introduction

In a recent paper [167], which extends earlier work [29, 46, 71], Sakaue *et al* used scaling arguments, based on the de Gennes polymer blob model [48], to predict when a polymer near the entrance to a narrow channel will be pushed into the channel by a flow field. They predicted that penetration will occur only above a threshold velocity flux

$$j_v = kT/\eta \tag{4.1}$$

where k is Boltzmann's constant, T is the temperature and η in the viscosity of the solvent. Surprisingly, as remarked in [167], there is no dependence in this formula on the length of the polymer nor on the dimensions of the channel. Physically this occurs because the penetration is a tunnelling phenomenon; there is a free energy barrier which the imposed flow must overcome to push the polymer into the constriction.

Recent developments in mesoscale simulation techniques mean that it is now feasible to perform convincing numerical investigations of the hydrodynamics of polymers moving within a flow field. In this paper we use one of these approaches, a coarse-grained representation of a polymer coupled to a lattice Boltzmann fluid [4, 5, 197], to investigate flow-driven injection into nanopores. We find that the Eq. (4.1) holds remarkably well, even for the short polymers and narrow channels that are accessible numerically.

The motion of polymers through constricted spaces has recently attracted considerable attention because of its relevance to, for example, biological applications of microfluidics [174] sequencing DNA by passing it through a nanopore [27, 218] and the passage of biomolecules through membrane channels [104]. The process by which a polymer moves through a channel can usefully be thought of in three stages. Firstly the polymer must find the opening, secondly it must enter the constriction, and only then can it move through the channel. Many authors have used simulations to model the final stage of this process. For example, there is work considering polymers driven through a pore by an electric field which addresses the role of charge[70, 132, 140], hydrodynamic correlations [68, 98], a crowded environment [74]

*published as: [127]: A.P. Markesteijn, O.B. Usta, I. Ali, A.C. Balazs, and J.M. Yeomans, Flow Injection of Polymers Into Nanopores, *Soft Matter*, 5:4575-4579, 2009

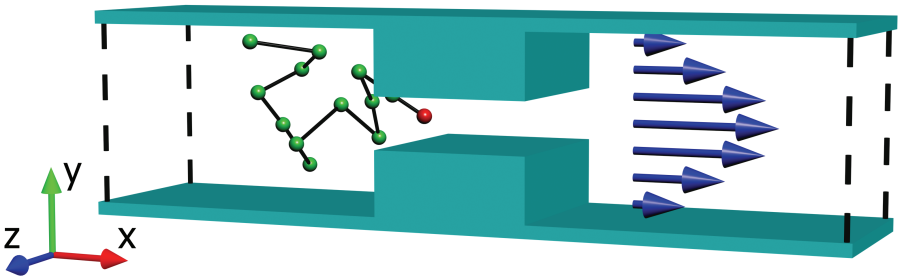

Figure 4.1: Geometry used in the simulations of a polymer chain translocating through a narrow slit. The size of the simulation box is (60,14,10). The dimensions of the slit are (12,h,10) and its left edge is at $x = 30$. All lengths are reported in simulation units.

and knots in the chain [93]. Polymers driven by a flow field have received less attention, but see [32, 182, 197] for examples concerning the behaviour of polymers moving in a Poiseuille flow field in a narrow channel.

In this paper we consider instead the second step of the translocation process. Rather than considering translocation times, we focus on a different question, namely, once the polymer has found the pore, will it be driven into it by the flow field or will it escape? This question was touched on by Matysiak *et al* [128], who used molecular dynamics simulations to model a polymer driven through a pore by an electric field. They found that the probability of translocation increased exponentially with the driving voltage at low voltages, which they interpreted as indicating an energy barrier, but became independent of the driving voltage at high voltages.

In Section 4.2 we describe our numerical algorithm. In Section 4.4 we present results for the dependence of the polymer translocation probability on the velocity flux, polymer length, temperature, solvent viscosity and channel width. The results are summarised in Section 4.4.5.

4.2 Numerical Method

To model the dynamics of a polymer chain entering a pore we use a coarse-grained, bead-spring model for the polymer chain and simulate the fluid by means of a fluctuating, lattice-Boltzmann method [113, 114, 115]. The beads (hydrodynamic centers) of the polymer and the lattice-Boltzmann model are linked by a frictional coupling [4, 5, 197]; this results in Oseen level hydrodynamics for the polymer chain. In this section we present the details of the algorithm.

4.2.1 Lattice Boltzmann Method

The lattice-Boltzmann model is based on a discrete analog of the continuous Boltzmann equation. The single particle velocity distribution function in the Boltzmann equation is replaced by a discrete version, $f_i(\mathbf{r}, t)$ which constitutes the fundamental quantity in the lattice-Boltzmann approach. $f_i(\mathbf{r}, t)$ describes the mass density of particles, at a lattice node $\mathbf{r}$ at time t, which have a discrete velocity $\mathbf{c}_i$ where i labels a lattice vector. It evolves according to a discrete Boltzmann equation [65]

$$f_i(\mathbf{r} + \mathbf{c}_i\Delta t, t + \Delta t) = f_i(\mathbf{r}, t) + \Delta_i\left[(\mathbf{r}, \mathbf{t})\right] \tag{4.2}$$

where Δ_i is the change in f_i due to collisions at the lattice nodes and Δt is the time step. The hydrodynamic fields, mass density ρ, momentum density $\mathbf{j}$, and momentum flux $\mathbf{\Pi}$, can be expressed as moments of the velocity distribution function

$$\rho = \textstyle\sum_i f_i, \quad \mathbf{j} = \rho\mathbf{u} = \sum_i f_i\mathbf{c}_i, \quad \mathbf{\Pi} = \sum_i f_i\mathbf{c}_i\mathbf{c}_i \tag{4.3}$$

where $\mathbf{u}$ is the macroscopic velocity at a particular node.

We use the standard 19 velocity model [65, 114, 115, 210], which consists of stationary particles and 18 velocities corresponding to the nearest ([100]) and next nearest ([110]) neighbor directions of a simple cubic lattice. The population density associated with each velocity has a weight a^{c_i} ($\sum_i a^{c_i} = 1$) that describes the fraction of particles with velocity $\mathbf{c}_i$ in a system at rest; these weights depend only on the speed $|\mathbf{c}_i|$. The optimum choice of weights for this specific lattice is

$$a^0 = \frac{1}{3}, \quad a^1 = \frac{1}{18}, \quad a^{\sqrt{2}} = \frac{1}{36}. \tag{4.4}$$

We describe the collision and the relaxation towards equilibrium using a 3-parameter collision operator [115] which allows for separate relaxation of the 5 shear modes, 1 bulk mode, and 9 kinetic modes. The post-collision distribution $f_i^\star = f_i + \Delta_i$ is written as

$$f_i^\star = a^{c_i}\left(\rho + \frac{\mathbf{j}\cdot\mathbf{c}_i}{c_s^2} + \frac{(\rho\mathbf{uu} + \mathbf{\Pi}^{neq,\star}) : (\mathbf{c}_i\mathbf{c}_i - c_s^2\mathbf{1})}{2c_s^4}\right) \tag{4.5}$$

where the sound speed $c_s = \Delta x/\sqrt{3}\Delta t$, and Δx is the lattice spacing. In previous studies [115, 197] it was found that the non-linear $\rho\mathbf{uu}$ term in Eq. (4.5) should be retained for suspensions of moving particles, since it maintains Galilean invariance and prevents an artificial cross-stream drift. Additionally, the domain of validity of Stokes flow is larger with the non-linear term in place.

The non-equilibrium momentum flux $\mathbf{\Pi}^{neq} = \sum_i f_i^{neq}\mathbf{c}_i\mathbf{c}_i$ relaxes due to collisions at the lattice nodes:

$$\mathbf{\Pi}^{neq,\star} = (1+\lambda)\overline{\mathbf{\Pi}}^{neq} + \frac{1}{3}(1+\lambda_v)(\mathbf{\Pi}^{neq} : \mathbf{1})\mathbf{1} \tag{4.6}$$

where $\mathbf{\Pi}^{neq} = \mathbf{\Pi} - \mathbf{\Pi}^{eq}$, $\mathbf{\Pi}^{eq} = \rho c_s^2 + \rho\mathbf{uu}$ is the equilibrium momentum flux, and $\overline{\mathbf{\Pi}}^{neq}$ indicates the traceless part of $\mathbf{\Pi}^{neq}$. The parameters λ and λ_v are eigenvalues of the linearized collision operator and are related to the fluid shear and bulk viscosities [115]

$$\eta = -\rho c_s^2\Delta t\left(\frac{1}{\lambda} + \frac{1}{2}\right), \quad \eta_v = -\frac{2\rho c_s^2}{3}\Delta t\left(\frac{1}{\lambda_v} + \frac{1}{2}\right). \tag{4.7}$$

The connection between the discrete evolution Eq. (4.2) and the Navier-Stokes equations can be established using a multi-scale expansion with the expansion parameter ϵ defined as the ratio of the lattice spacing and a macroscopic length scale. One can show that in the hydrodynamic limit ($\epsilon \ll 1$) the lattice-Boltzmann equation leads to the Navier-Stokes equations with corrections of order u^2 and ϵ^2 [65, 210].

A pressure-driven flow can be modeled using the lattice Boltzmann algorithm by introducing an external force density $\mathbf{F}^{ext}$. This is done by modifying the evolution equation to

$$f_i(\mathbf{r} + \mathbf{c}_i\Delta t, t + \Delta t) = f_i^\star(\mathbf{r}, t) + F_i^{ext}(\mathbf{r}, t). \tag{4.8}$$

In the presence of external forcing, more accurate solutions to the velocity field are obtained if it is defined to include a portion of the momentum added to each node by the force [115]

$$\mathbf{j}' = \rho\mathbf{u}' = \sum_i f_i\mathbf{c}_i + \frac{1}{2}\mathbf{F}^{ext}\Delta t. \tag{4.9}$$

In its original form the lattice-Boltzmann approach does not lead to fluctuating hydrodynamics. However, it is possible to introduce fluctuations by including random terms in the non-equilibrium stresses during the collision process. We follow the prescription of [113, 114, 115] which satisfies a discrete fluctuation-dissipation theorem. The magnitude of the fluctuations can be controlled in a convenient way by using an independent variable, T, identified as the temperature.

To introduce solid boundaries in the system we use a bounce-back collision rule [65] in which incoming fluid particles are reflected back towards the nodes from which they originated. This results in no-slip boundary conditions.

4.2.2 Polymer Chains

We model the polymer chain as a series of N beads of mass m connected by Hookean springs with potential

$$U_S(r) = \tilde{k}(r-b)^2 \tag{4.10}$$

where the equilibrium length of a bond b is chosen equal to the lattice spacing Δx. The stiffness $\tilde{k} = 300$, ensures low fluctuations in the bond length compared to the radius of gyration of the polymer chain. Excluded-volume interactions between beads are modeled using a truncated DLVO-like potential in the limit of a point particle,

$$U_E(r) = U_0 \frac{\exp(-\kappa r)}{r}, \tag{4.11}$$

with DLVO parameter $U_0 = 0.1$ and the inverse Debye-Huckel screening length $\kappa = 80$.

The equation of motion of the polymer beads are solved in an inertial form:

$$m\frac{d^2\mathbf{X}_i}{dt^2} = -\boldsymbol{\nabla}_i U(\mathbf{X}^N) + \mathbf{F}_i^f \tag{4.12}$$

where $\mathbf{X}_i$ is the position of bead i, U is the total potential energy, which depends on the positions $\mathbf{X}^N$ of all the beads in the chain, and $\mathbf{F}_i^f$ is the hydrodynamic force exerted on bead i by the fluid. It was shown that the large time-scale separation between the dynamics of the polymer and the individual monomers allows time for the hydrodynamic interactions to reach a quasi-steady state, without imposing non-inertial equations of motion at each time step [197]. We integrate Eq. (4.12) using a first order implicit scheme. The evaluation of the forces arising from the potential, U, can be carried out using time-steps shorter than Δt.

The beads of the polymer chain are coupled to the lattice Boltzmann fluid through a frictional force that depends on the relative velocity of the bead and the fluid

$$\mathbf{F}_i^f = -\xi_0 \left[\dot{\mathbf{X}}_i(t) - \mathbf{u}(\mathbf{X}_i, t)\right] + \mathbf{F}_i^r. \tag{4.13}$$

where ξ_0 is a friction coefficient which is computed using the Stokes drag, $6\pi\eta r_a$, where r_a, the hydrodynamic radius, is set to 0.32 for all simulations. Since the beads move continuously whereas the fluid field is discrete in space, the fluid velocity is interpolated to the center of mass position of each bead using a trilinear interpolation scheme [4, 5]. The random force $\mathbf{F}^r$ serves to balance the additional dissipation caused by using a frictional coupling instead of a no-slip boundary condition on the bead surfaces [5]. Unlike for Brownian dynamics routines, here $\mathbf{F}_i^r$ has a local covariance matrix since the fluid satisfies its own fluctuation-dissipation theorem [5, 197]

$$\langle \mathbf{F}_i^r(t)\mathbf{F}_i^r(t')\rangle = 2kT\xi_0\delta(t-t')\mathbf{1}. \tag{4.14}$$

In order to conserve local and global momentum, the accumulated force on a bead is distributed back to the fluid using the same trilinear interpolation [4, 5]; this completes the coupling between the bead and the fluid. The approach described leads to Oseen-level hydrodynamic interactions between the beads [197].

4.3 Simulation Details

We simulate the motion of a single polymer chain in a box of dimensions (60,14,10), where all lengths are reported in lattice units Δx. In the middle of this geometry we create a narrow slit of size (12,h,10) as depicted in Fig. 4.1. The slit lies between $x = 30$ and $x = 42$. Periodic boundary conditions are imposed along $\hat{\mathbf{x}}$ and $\hat{\mathbf{z}}$, and no-slip boundary conditions on the other surfaces. The slit geometry used here is expected to give the same scaling as the pore geometry used to derive Eq. (4.1) if the motion of the polymer into the slit is essentially one-dimensional. This was observed to be the case; eg the probability of hairpins moving into the channel was negligible.

The polymer chain comprises N beads connected by $N-1$ springs. Initially the first bead of the chain is pinned just inside the left-hand entrance of the slit (at $x = 30.9$) and the rest of the beads are allowed to equilibrate for 50000 time steps. The beads are prevented from entering the slit during the equilibration by means of a soft virtual wall. The wall is then

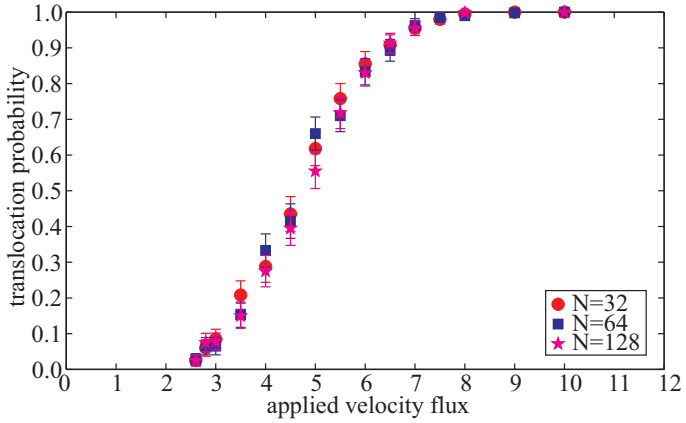

Figure 4.2: Probability of translocation through the slit as a function of applied velocity flux for different lengths of the polymer chain.

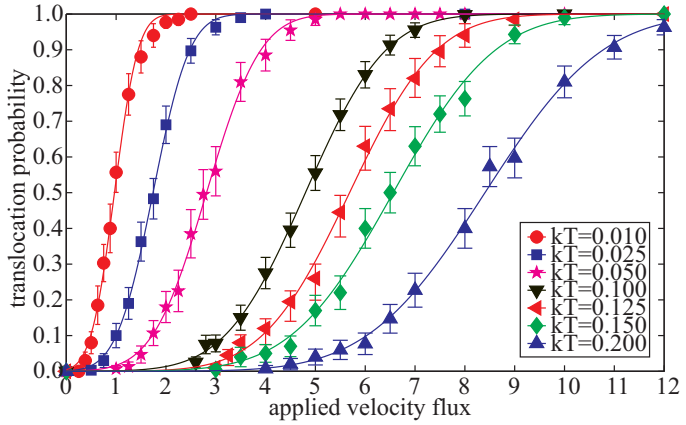

Figure 4.3: Translocation probability as a function of applied velocity flux for different temperatures. The continuous lines indicate fits of the data to functions of the form of Eq. (4.16).

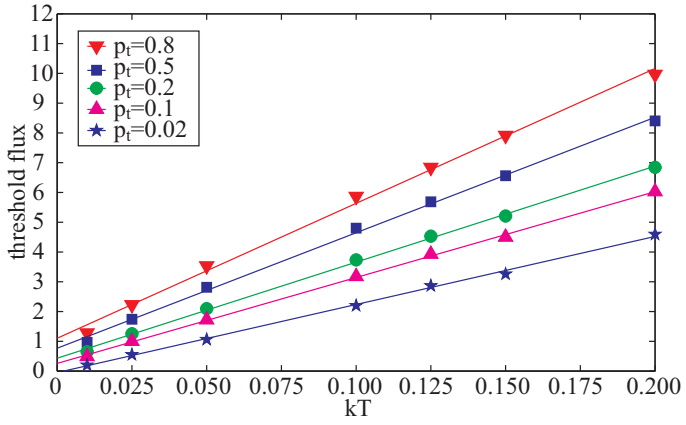

Figure 4.4: Applied velocity fluxes which lead to a probability of translocation p_t, as a function of temperature.

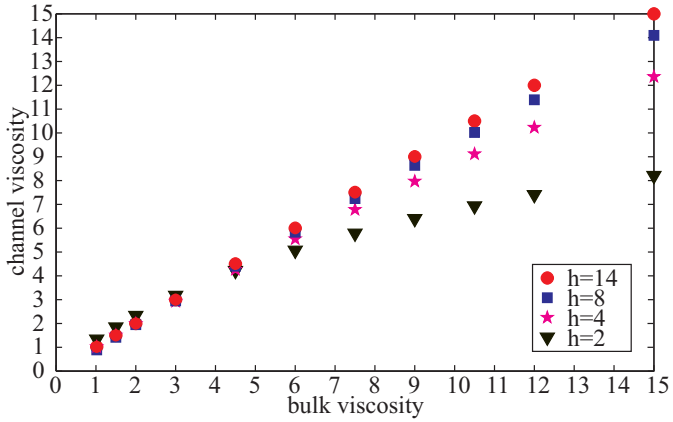

Figure 4.5: Comparison of the measured viscosity in the channel to the viscosity in the bulk for different channel widths h. The values have converged by $h = 14$.

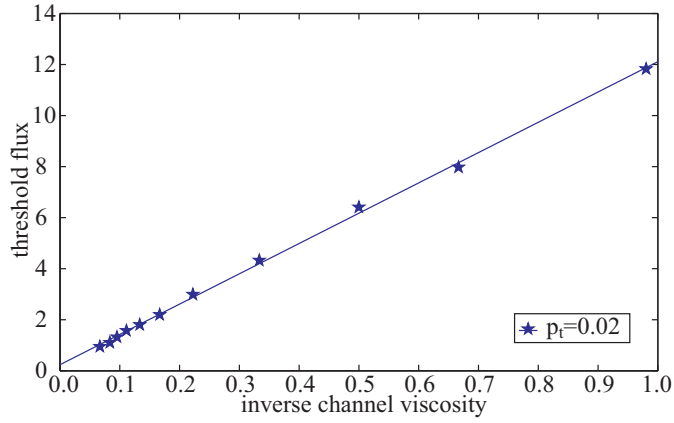

Figure 4.6: Threshold flux, which leads to a probability of translocation $p_t = 0.02$, as a function of the inverse channel viscosity. The line is a linear fit to the data.

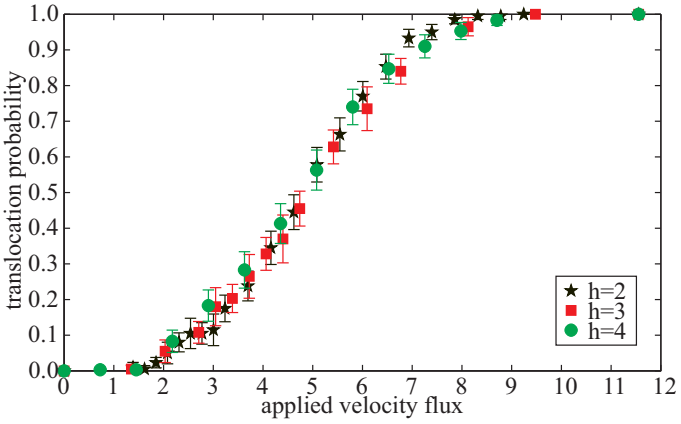

Figure 4.7: Translocation probability as a function of applied velocity flux for different channel widths h.

removed and a pressure driven flow is applied. Assuming two dimensional Poiseuille flow within the slit, the mean velocity flux per unit width of the slit is

$$j_v = h^3 F^{ext}/12\eta. \tag{4.15}$$

Fluxes are measured with respect to a reference value $j_v^0 = 1.1111 \times 10^{-6}$, which corresponds to $\lambda = 1$, $\rho = 36$, giving $\eta = 6$, $h = 2$ and $F^{ext} = 10^{-5}$. The simulation is stopped when all the beads are out of the slit, on either side. A successful event, translocation, corresponds to the polymer ending up on the right hand side of the geometry, $x > 42$. We repeat the simulations at different temperatures, kT, for chain lengths, $N = 32, 64, 128$, and different viscosities, η and slit heights h, varying the applied flux. For each parameter set we perform 300 simulations to calculate the probability of translocation, P.

4.4 Results

4.4.1 Dependency on the Number of Beads "N"

Fig. 4.2 shows the translocation probability, plotted as a function of the applied velocity flux, for three different values of N, the length of the polymer chain. Other variables are kept constant at $kT = 0.1$, $\eta = 6$ and $h = 2$. The error is obtained by calculating the standard error of the mean, which is: $S_E = {}^{s}/\sqrt{n}$, where s is the standard deviation of the results of the simulation (0 means failed, 1 means success) and n are the number of simulations ($n = 300$). The errorbars in the figure indicate the 95% confidence limit of the sample mean, which is $1.96 \times S_E$. As the flux is increased, the translocation probability crosses over smoothly from 0 to 1. Fig. 4.2 shows clearly that the probability of translocation is independent of the length of the chain, in agreement with the prediction of Eq. (4.1) [167].

4.4.2 Dependency on the Temperature "kT"

Next we consider Fig. 4.3, which shows the probability of translocation as a function of applied velocity flux for a range of values of the temperature kT. The polymer length in these simulations was held fixed at 128 and each data point results from 300 runs. The critical flux needed to push the polymer through the constriction increases with temperature, as the entropic penalty of confinement becomes greater. Moreover, the distribution becomes broader ($\sigma^2 \sim kT$). This is because the probability that the polymer fluctuates to escape the constriction, even when translocation is possible, is greater with increased temperatures and therefore it becomes more difficult to achieve a translocation rate of unity.

To define a critical flux more quantitatively it is helpful to fit the translocation probability curves to a continuous function. A choice that fits the data well is

$$p(j_v) = \frac{1}{2}\left\{1 + \mathrm{erf}\left(\frac{j_v - j_v(\frac{1}{2})}{\sqrt{2}\sigma}\right)\right\} \tag{4.16}$$

where we define $j_v(p_t)$ is the value of the velocity flux for which a threshold fraction p_t of the polymer chains translocate. The variation of j_v with temperature, for different values of p_t is plotted in Fig. 4.4. The plotted fluxes shows a linear dependence on temperature [167]. This linearity is independent of the definition of the threshold probability. The slope and intercept of the lines do, however, show a weak dependence on the choice of threshold. Henceforth we shall present results for $p_t = 0.02$ and we shall term the $j_v(0.02)$ as the threshold flux. A reason for choosing this value of p_t is that it is least sensitive to the chain's ability to fluctuate away from the tube before it is pushed through. Physically this is reflected in the threshold flux tending to zero as $kT \to 0$, as expected.

4.4.3 Dependency on the Viscosity "μ"

We now consider the dependence of the threshold velocity flux on the fluid viscosity. For a non-confined fluid the viscosity is determined by the value of λ and ρ through Eq. (4.7). However, the value of the velocity and therefore also the apparent viscosity within such a small channel is strongly affected by the details of the boundary conditions and number of

discrete points inside the small channel. In other words, from a computational view, the effective viscosity of the fluid inside the channel is not necessarily the same as the viscosity in the bulk. Therefore the effective viscosity of the fluid inside the channel was measured by running a simulation without the polymer and fitting the velocity profile within the slit to the formula for Poiseuille flow. Fig. 4.5 plots the relationship between viscosity in the bulk and in the channel as a function of the slit width h. The two values have converged by $h = 14$.

Results for the translocation probability as a function of the applied velocity flux, analogous to those presented in Fig. 4.3, were obtained for a chain of length $N = 128$ for several different values of the channel viscosity. From these we extracted the threshold velocity flux $j_v(0.02)$. This is plotted as a function of the inverse channel viscosity in Fig. 4.6. The data shows good agreement with the dependence $j_v \sim T^{-1}$ predicted by Eq. (4.1). The results are likely to be affected by, for example, the difficulty in measuring the viscosity precisely near the entrance to the channel.

4.4.4 Dependency on the Slit Width "h"

As a final check on Eq. (4.1) we repeated the simulations for different values of the slit width h. We considered chains of length 128 and each data point was the result of 400 runs. The results, summarised in Fig. 4.7, show that the translocation probability for a given value of the applied velocity flux, is indeed independent of h, at least for the narrow slits it was feasible to simulate.

4.4.5 Conclusions

The numerical results presented in this article provide a striking confirmation of Eq. (4.1) which states that the threshold velocity flux needed to capture a polymer within a narrow constriction is independent of the length of the polymer, depends linearly on the temperature, and is inversely proportional to the fluid viscosity. Eq. (4.1) follows from a scaling argument, based on the de Gennes blob model of polymers, and hence is expected to be applicable for long chains: however, we find that the expression can also be used to describe short chains in a channel which has dimensions on the order of the Kuhn length of the polymer. It would be interesting in future work to determine the extent to which the threshold velocity flux depends on topological and chemical details of the polymer chain and on its interactions with the channel walls. Numerical checks of other predictions of the blob model, for example its application to electrophoresis [166], would also be of interest.

Acknowledgements
We thank T. Sakaue for helpful discussions. A.P. Markesteijn was supported by the MicroNed program and the EU INFLUS project. J. M. Yeomans acknowledges support from the Office of Naval Research.

Chapter 5

Conclusions and Recommendations

In this final chapter the main conclusions that are presented in the previous chapters will be summarised and discussed. This is done in Section 5.1. Furthermore, in Section 5.2 several possible directions for future research with respect to the topic of this thesis will be suggested.

5.1 Conclusions

The essence of the coupled simulation is to accurately model phenomenon, like wall effects, that basically happen on the molecular scale (possibly non-continuum), while this same phenomenon can have a large effect on the macroscopic variables, like velocity, temperature, or motion of nano-sized or micro-sized beads/particles. The coupling is necessary, because an attempt to model both the microscopic and macroscopic world with molecular dynamics alone, would be impossible because of computational limits. However, molecular dynamics can successfully be applied to systems that have a size of up to $\mathcal{O}\left(10^{-8}\right)$ metres. On the other hand, simulations with only continuum methods result in assumptions of boundary conditions, but do produce reasonably accurate solutions where the typical size of the device that is studied is on the order $\mathcal{O}\left(10^{-5}\right)$ metres or larger. Therefore, the advantage of the coupled simulation is that systems that are between the two orders, $\mathcal{O}\left(10^{-8}\right)$-$\mathcal{O}\left(10^{-5}\right)$ metres, can be simulated effectively. The coupled simulation will also take away these assumptions of the boundary conditions, where they are now replaced with a better representation of what is happening.

The main goal of the research presented in this thesis was to develop a coupled numerical simulation for dense liquids where one domain is represented with molecular dynamics, while the other domain is treated by the conventional continuum equations. The two domains are coupled using the Schwarz alternating method, where the two domains are solved separately and are coupled in a so-called overlap region. Inside the overlap region the two domains must solve for the same solution and at either end of the overlap region, the communication between MD and the continuum takes place. It is the correct communication between the two domains that enables a successful coupling. This way it is also possible to not only couple the macroscopic variable of velocity, but also other variables like temperature or stress, either separately or combined. However, before this can be accomplished, it is necessary to obtain more information about where and when a coupling is possible and under which assumptions these coupled simulations give a correct solution. Therefore, the present work also demonstrated several MD simulations/studies to investigate the possibilities and limitations of MD, while also the limitations of pure continuum methods in similar conditions are investigated. Further, previous research mainly demonstrated how coupling can be achieved and illustrated the method by doing simulations where the liquid consisted of

simple (noble) atoms, i.e. no molecules. Here, it also shown how a simulation with molecules, especially water, can also be used in a coupled approach and results of this coupled simulations are shown.

In Chapter 2 the foundations of the numerical models for continuum mechanics and molecular dynamics are explained and they are compared to each other. The main results are:

Large deviations between continuum mechanics and MD are especially noticeable near the solid walls of nano-sized channels or near obstacles and are local.

- They are the result of the interaction of the atoms in the liquid with the atoms in the solid wall and can also be observed experimentally.
- They are *not* a result of an applied flow or other external source.
- The sampled profile of macroscopic variables, especially the density profile, shows large peaks and valleys in the vicinity of the wall and only far away from the wall the macroscopic variables show their (expected) continuum value without variations.

Although these large variations in the sampled macroscopic variables indicate large non-continuum effects, even for small nano channels of about 5 nm, near continuum-like behaviour can be extracted from the results.

- It is shown how a Poiseuille flow can be generated and how the velocity profile can be used to curve fit a Poiseuille velocity profile. This is done to determine the viscosity of the liquid simulated.
- The simulation using argon showed that the value of the viscosity is $\mu \approx 2.0$.
- For water the value of viscosity for four frequently used water models are presented and compared to the experimental value. The results showed that the TIP4P/2005 model has outstanding accuracy (5% to 10% deviation), while both the TIP4P/EW and SPC/E model showed reasonable accuracy (15% to 30% deviation), and the TIP4P model performed the worst with deviations up to 60%.
- These results are very important for the coupled simulation. Unlike what is the case for MD, where the viscosity is a result of the simulation, the continuum flow equation requires a value for the viscosity in order to solve the equation. The coupled solution will diverge as a results of the mismatch between the two values of viscosity.

In general, the results of the simulations showed that in a nano channel with a height of about 8 nm yield very good overall continuum-like behaviour. However, this does not yet mean that a pure continuum method to compute the flow inside this channel is advisable.

- The variations in the values of several properties are noticeable until about 2.6 nm distance away from *each* wall of the channel. For the nano channel with a height of 8 nm this means that 65% of the nano channel experiences non-continuum effects, while for a nano channel with a height of approximately 58 nm there are still 10% non-continuum effects present.

There are also different reasons why not to use a continuum method to simulation certain phenomena, because MD simulations do have some unique benefits and allow for more controllability:

- In an MD simulation the wall-fluid interaction is directly simulated, which gives the possibility to model actual roughness or a small contamination present inside the wall, and the effect this has on the resulting flow. It also means that the interaction strength between the wall and the fluid can be changed. The results of such a simulations showed that the average velocity profile inside a nano channel with roughness resembles the expected continuum behaviour, i.e. roughness causes a lower flow through the channel if the same force is applied. However, the MD simulation allowed much more detail to be implemented in the simulation. A simulation with a reduced wall-fluid interaction showed how the obtained velocity profile in the nano channel shows a shift of the velocity, i.e. an apparent wall slip occurs, equivalent to the specification or assumption of a certain slip velocity in the case of a pure continuum computation.

- Furthermore, with MD several other phenomena can be simulated that are difficult or even impossible with a continuum technique, like the nano-jet and nano-jet breakup, which has application in future inkjet printing. Besides the fact that MD naturally handles the actual break-up of the jet, the added benefit of a MD simulation is the possibility to accurately model wall-fluid interactions like wetting, hydrophilic or hydrophobic behaviour at the nozzle of the inkjet.

However, with MD simulations care must be taken, because frequently used values of the cutoff radius, which range from 2.3σ ***to*** 4.0σ***, are too low to accurately model several important phenomena, especially surface tension or droplet break-up.***

- This is demonstrated by a simulation of a nano-sized inkjet printer and the break-up of a nano-jet with different cut-off radii and a relatively new method, where a long-range technique is used to compute the long-range attraction without cut-off, while the other interactions are computed with a very small cut-off radius.
- The long-range technique for the intermolecular interactions showed similar results compared to the simulation that employed a very large cut-off radius, while the computational effort is less.
- The employed cutoff radius did have a big impact on the computed droplet creation process, e.g. the elongation of jet and ultimate shape and size of the droplet.
- The wall-fluid interaction had a major impact on both the droplet size and velocity.
- Surprisingly, the cutoff radius did not seem to have a lot of effect on the computed droplet velocity, which is an important model parameter for inkjet printers.

Finally, by comparing continuum results with experimental result, it is shown that continuum techniques, especially the ones describing electrokinetic effects, are reasonably accurate enough for a nano-sized device where the height is only 150 nm.

- The computation was accomplished by taking into account the specific large aspect ratio between the height and the other dimensions of the device.

In Chapter 3 it is explained how a domain where MD is employed and a domain where any continuum technique is employed, can be coupled. This can be accomplished by using the Schwarz alternating method, where the two domains are coupled inside an overlap region. ***The main task involved in the coupling are the correct specification of the boundary conditions on the MD and continuum domain.***

- The extraction of the correct boundary condition for the continuum domain from the MD domain is a straightforward process, because it only involves sampling of statistics from the MD domain.
- Imposing the correct boundary conditions on the MD domain which are obtained from the continuum domain is more challenging and three problems needed to be solved:
    1. *Imposing the desired macroscopic variable to the MD domain.* This can be accomplished by introducing an additional force inside the overlap region, where the coupled results can be improved by adjusting this force to take the (continuum) gradient into account inside the overlap region.
    2. *Controlling the mass flux and consequently the atoms or molecules that now can leave and enter the MD domain.* A very simple but elegant solution is to use a simple statistical probability, where the likelihood of one atom to bounce back or to be reinserted is based on the sign of a certain counter. The whole insertion criterion can be controlled by the relative weight that are given to the counter when an atom is bounced back or reinserted.
    3. *Solving the problem that near the MD domain boundary incomplete intermolecular interactions take place.* The incomplete interaction is a results of the fact that the continuum domain does not simulate atoms or molecules by definition, while molecular dynamics requires the individual position of the atoms in order to compute the interaction. The problem is that the missing interactions

result in unnatural variations in sampled variables near the MD boundary. Previous research focused on compensating the missing information by providing the molecular dynamics simulation meaningful information through a so-called boundary force function. In this work, an adaptation of this research is given. If the correct boundary force function is used, it does eliminate the unnatural variations in the sampled macroscopic variables near the MD boundary.

The same principles behind the coupled simulations using a liquid that consisted of single non-polar atoms, i.e. argon, can also be used for coupled simulation using water. However, before a successful coupling could be achieved, especially the problem of incomplete interactions near the MD boundary needed to be solved. The main results are:

- Previous research on the boundary force function resulted in a single function that is applied to the centre of mass of each water molecule. Although this function is able to eliminate the density variations near the boundary, it is not successful in preventing incorrect orientation of the water molecule near the boundary.
- In this work is was investigated whether a boundary force function for each separate atom inside the water molecule and for each type of interaction, i.e. short and long ranged interactions, would result in correct orientations near the boundary. However, the results of the simulations revealed that doing so, does not influence the orientation of the water molecule near the MD boundary.
- On the other hand, the boundary force function obtained in this work is very different to the boundary force function obtained by others, while the variations in the sampled macroscopic variables are still eliminated. This result indicates that, unlike in the case of argon, no unique (single) boundary force function can be found for water as such.

To demonstrate the applicability of the coupled approach, several coupled simulation were performed. The main results of these simulations are:

The convergence of the Schwarz alternating method, used to couple two continuum domains and one MD domain, behaves exactly the same as when the method is used to couple two continuum domains.

- The relative error decreases exponentially with the number of iterations and the needed number of iterations to obtain a converged solution scales inversely proportional to the relative overlap size.

The coupled simulation can be seen as a new boundary condition for the continuum, where the value is now more accurately supplied by the communication of the MD and continuum domain. This was demonstrated by simulating Poiseuille flow of argon and water inside a large nano channel. The main results of these coupled simulations are:

- The non-continuum effects near the wall are simulated accurately by MD and no expensive MD computation time is wasted on the part that resembles a continuum.
- The total solution is the same as the solution obtained with a full MD simulation, while the total computational effort is reduced by a factor 4 for the simulation of a nano channel with a height of about 23 nm and this factor improves if the channel height is increased, while there is virtually no drop in accuracy compared to the full MD simulations.
- This type of simulation can be used for cases where the wall-fluid interaction is very important and cannot be correctly simulated by standard continuum techniques, like phenomena such as wetting and velocity slip in microchannels.
- For the successful coupling it is important to match the value of the viscosity. A mismatch of the viscosity results in a velocity profile that is stretched and non-parabolic compared to the expected velocity profile.
- It is also important to include the continuum (velocity) gradient inside the overlap region in the MD domain. If no gradient is included the velocity profile shows a plateau or flattening, which also influences the local shear and the surrounding velocity profile and therefore the communication between the MD and continuum domain. This results in a coupled solution that is shifted with respect to the expected solution.

The coupling of MD and continuum also enabled the specification of non-periodic boundary conditions for MD systems, which are difficult or impossible to implement in a pure MD case. This was demonstrated by a two-dimensional coupled simulation of a nanowire inside an uniform flow of argon. The main benefits and results of this type of simulation are:

- It investigates the influence on the flow of *one* nano-sized obstacle, instead of the one and all its periodic images.
- The nano-sized obstacle can be any molecular structure of interest, e.g. a nanotube, nanowire, or a complex molecule like DNA, while a pure continuum technique has to use many assumptions.
- The resulting velocity profiles showed that there is a difference in solution between the coupled situation and the pure continuum situation, especially noticeable near the nanowire. This is a direct effect of the interatomic interactions between the fluid and the nanowire, which show density variations near obstacles.
- The density variations also mean that the local viscosity is affected. The value of the viscosity near the nanowire is much higher than the value of viscosity in bulk conditions.

The principles behind the coupling of domains can also be applied to other macroscopic variables than velocity, for example temperature. In this work a qualitatively study was performed on a single nano-sized particle inside a temperature gradient field by coupling the MD domain and the continuum domain, effectively investigating thermophoresis in liquids. The main results are:

- The total force on the particle is a linear function with respect to the applied temperature gradient.
- The comparison between two different sized particles showed only a small difference between the forces when they experience the same temperature gradient.
- In all cases the force had an opposite sign than the temperature gradient, indicating that in all cases studied, the particle wanted to migrate towards the cold side of the domain. This conclusion was also drawn by others, as this seems to be case with all simulations that use the Lennard-Jones 12-6 potential applied to argon atoms.
- The results from a simulation that employed a larger cut-off radius showed that the total force on the particle changed by as much as 20% when the cut-off radius was increased from $r_c = 2.5\sigma$ to $r_c = 4.0\sigma$. This again indicates that the commonly accepted cut-off radius of 2.5σ is too small for these kind of applications.

In Chapter 4 a different type of coupling is presented, where a mesoscale simulation is used to measure the strength of the velocity flux needed to push a polymer into a narrow channel. Here, the dynamics of the polymer chain is modelled using a coarse-grained, bead-spring model for the polymer chain, while the fluid is simulated by means of a fluctuating, lattice-Boltzmann method. The beads (hydrodynamic centers) of the polymer and the lattice-Boltzmann model are linked by a frictional coupling, which results in Oseen level hydrodynamics for the polymer chain. Using this technique it was investigated whether the prediction, based on the de Gennes polymer blob model, that penetration will occur only above a threshold velocity flux, is valid. Furthermore, the threshold velocity flux should only depend linearly on the temperature and inversely proportional to the fluid viscosity. The simulation results showed excellent agreement with the prediction, even for the short polymers and narrow channels that are accessible numerically. It was also shown how the threshold velocity flux is independent of the length of the polymer chain or the height of the channel.

5.2 Recommendations

The recommendations for further research are formulated as follows:

Investigating coupled simulations with MD and continuum for microfluidics.

- In this thesis, coupled simulation of mostly nanosized channels were investigated. Here the atomic details near the wall of the channel are simulated with MD, while the remainder of the channel is modelled with a continuum method. The coupling approach discussed in this thesis can also be used to simulate a microfluidic device, e.g. a channel that is 5 micrometre in height. However, the problem is collecting enough statistics. For example, to keep the velocity to acceptable limits inside the whole micro channel in the case of Poiseuille flow, would mean that the value of the velocity inside the overlap region used for the coupling of the MD domain and the continuum domain is nearly zero. This means that the velocity profile will be very hard to sample, i.e. "bad" statistics involved with the very low velocity compared to the natural vibrations of the atoms. A possibility to improve the statistics is to either sample more atoms, i.e. simulating a larger MD domain, or collect the statistics over a greater length of time. In both cases, the computational effort will increase more or less linearly with the number of atoms or time steps and will soon become too cumbersome for the average computer system. This also leads to the next recommendation. However, note that any relatively easy sampled property, like the density or the temperature, can still provide very useful information for the coupled simulation results.

Improving the statistics by using different sampling methods or taking advantage of increasing multiprocessing and GPU computing.

- Molecular Dynamics is a powerful computational technique to simulate the behaviour of liquids without almost any need for modelling, but the method is very computationally expensive. The number of atoms used in the coupled simulations shown in this thesis were on the order of $\mathcal{O}\left(10^4\right)$ and are sampled over the number of time steps on the order of $\mathcal{O}\left(10^6\right)$. These simulations consume several hours to several days of CPU power on one processor. However, reasonable results could only be obtained within this time if the sampled macroscopic variable, especially velocity, was significant compared to the natural fluctuations present in any MD computation. For the velocity this means that simulated values are on the order of $\mathcal{O}\left(10^4\right) - \mathcal{O}\left(10^6\right)$ larger than commonly encountered values in experiments, which is not a problem as long as the expected behaviour scales accordingly. However, MD is an ideal target for multiprocessing and therefore can take advantages of the current trends of the use of multi-core CPUs and GPU computing. Doing so, could mean statistics could be improved by simulating more atoms in the same time, therefore decreasing the noise-ratio of the sampled macroscopic variables. On the other hand, by implementing specialised time-domain decomposition techniques, the time integration could also be distributed amongst many processors and again the statistics could be improved.

Extending the method used to eliminate the unnatural variations in sampled variables near the MD/Continuum boundary.

- It is shown how the so-called boundary force function can be used to eliminate the density variations as a result of the incomplete molecular interactions near the MD/Continuum boundary. For the case of a single atom liquid, like argon, this function is unique and is able to eliminate the variations completely. However, for a molecule, that has a certain position and rotation associated to it, this is not necessarily the case. Further research should therefore investigate how the boundary force function can be improved to include more parameters than just the unnatural variations near the boundary or how a different approach could be used to do the same.

Further improvement of the coupling of continuum gradients.

- In this thesis it is shown how the inclusion of the first (continuum) gradient inside the overlap region can improve the coupled results considerably. The implementation of

the velocity gradient is simply done by adjusting the force inside the overlap-region, while the temperature gradient is applied through careful thermostatting. However, further research should also include higher gradients. A possible start point could be to use the characteristics of the continuum method employed. For example, in the cases reported in this work, the continuum method was predominately the Finite Element Method, and therefore the inclusion of the gradient could use the finite element shape functions to couple these to the additional MD force.

Improving the coupling of long-range interactions.

- Although long-ranged interactions are very efficiently implemented into the MD domain, in the present study they are not satisfactorily implemented in the coupled simulations. More research should be conducted in this area. Possible start points for the long-ranged electric effect could be the direct coupling of the PPPM charges (which can also be used to model the long-range term in the Lennard-Jones potential, i.e. the r^{-6} term) to the continuum Poisson equation.

Extending the coupling approach to include a MD domain and/or continuum domain that is allowed to move.

- The simulations shown in this thesis concentrated on the flow round, or phenomenon near, a static object. In other words, the MD domain remained at the same position inside the continuum domain and did not change size during time. Therefore, to prevent the particle of interest to move outside the static domain, the particle was pinned to its position. For the realistic simulation of thermophoretic motion, electrokinetic flow, or any situation where the part of the domain that cannot be simulated with a continuum method, this part should have the freedom to move. One intermediate solution to solve this problem is to still keep the particle inside the centre of the MD domain, but at the same time move the entire MD domain (and its atoms) with the amount that the particle should have moved when it would not be pinned. However, this also means that the communication of the boundary conditions between the MD domain and continuum domain must change accordingly. Therefore, further research should point out whether these kind of simulations can still be performed efficiently using the Schwarz alternating method.

Investigating the extend of the prediction of the critical velocity flux for translocation based on a de Gennes blob model of the polymer.

- It would be interesting in future work to determine the extent to which the threshold velocity flux depends on topological and chemical details of the polymer chain and on its interactions with the channel walls. Numerical checks of other predictions of the blob model, for example its application to electrophoresis, would also be of interest.

Appendix A

Electroosmosis

In this section the governing equation for electroosmosis are given. It will be shown that although the general theory is well understood, in most cases (mostly microfluidics) very good approximation can be made. This is primarily because of the great difference in typical length scales of the electric double layer and the device itself. Only in the case of nanofluidics there might be a need to use the full theory, however it will also be shown that in the extreme case, where the device is much smaller than the electric double layer, again good approximations can be made.

A.1 General Theory / Poisson-Boltzmann equation

The fundamental law of electricity (and magnetism) are Maxwell's equations. The magnetic effects can be proven to be negligible small compared to the electric effects in almost all micro- and nanofluidic applications, and in the case of (DC) electroosmosis the electric field is also static. In this case Maxwell's equations reduce to the *Poisson equation*:

$$\nabla \cdot \left(\epsilon \vec{E}\right) = \frac{\rho_e}{\epsilon_0} \tag{A.1}$$

where $\vec{E}$ is the electric field, ϵ the relative permittivity of the electrolyte solution, ϵ_0 the permittivity of free space, and ρ_e the net electric charge density. Another result from Maxwell's equations is that the electric field can be derived from a *electric potential* Ψ as $\vec{E} = -\nabla\Psi$. The *net electric charge density* in this case is due to all ions present per unit volume, which is

$$\rho_e = \sum_i c_i z_i e \tag{A.2}$$

where c_i is the concentration of species i, z_i the valency of species i, and e the electron charge.

The fundamental law of the transport of ions in a (dilute) solution is the *Nernst-Planck equation*, which sums three types of transport. The first term describes the movement of ions based on concentration gradients, which is Fick's law of diffusion. The second term describes the movement due to electric fields. The third and last term is the particle advection term. The total particle flux for species i, $\vec{J}_i$, can now be written as

$$\vec{J}_i = -D_i \nabla c_i - v_i c_i \nabla\Psi + c_i \vec{u} \tag{A.3}$$

where D_i the diffusion constant for species i according to Fick's first law, v_i the mobility of species i, and $\vec{u}$ the velocity vector. Equation A.3 can be written in a more convenient form when the Nernst-Einstein equation, which gives a relation between the diffusion constant and the mobility, and the continuity equation for species i, is used. The resulting equation is

$$\frac{\partial c_i}{\partial t} + \vec{u} \cdot \nabla c_i = D_i \nabla \cdot \left(\frac{z_i e}{k_b T} c_i \nabla\Psi\right) + D_i \nabla^2 c_i \tag{A.4}$$

where k_b is the Boltzmann constant, and T the absolute temperature.

The fundamental laws for the flow of (Newtonian) fluids, are the equation of continuity and the Navier-Stokes equation. However, the Navier-Stokes equation needs to be modified such that the interaction between the ions and the applied electrical field are taken into account. This can be accomplished by adding an extra body force term. This extra body force term is derived from the Maxwell electromechanical stress tensor [117] for a body in a electric field by using equation 2.4 and equals:

$$\vec{f_e} = \epsilon_0 \nabla \cdot \left(\epsilon \vec{E}\right) \vec{E} - \frac{\epsilon_0}{2} \vec{E} \cdot \vec{E} \nabla \epsilon + \nabla \left[\frac{1}{2} \epsilon_0 \vec{E} \cdot \vec{E} \rho_f \left. \frac{\partial \epsilon}{\partial \rho_f} \right|_T \right] \tag{A.5}$$

where ρ_f is the fluid density. The first term is the electrostatic force density. The second term represents a force which appears whenever an inhomogeneous dielectric is placed in an electric field respectively. This term is the driving mechanism for dielectropheresis. The last term is the electrostriction, which is the pressure due to electric effects. For typical applications of electroosmotic flow only the first term is required. Also, please note that this term can be rewritten with aid of equation A.1. The final governing equations for the flow are the continuity equation, equation 2.1 and the *modified Navier-Stokes equation:*

$$\rho_f \left(\frac{\partial \vec{u}}{\partial t} + \left(\vec{u} \cdot \nabla\right) \vec{u} \right) = -\nabla p + \mu \nabla^2 \vec{u} + \rho_e \vec{E} \tag{A.6}$$

Therefore the general theory that describes electroosmotic flow is a system of coupled differential equations, namely equations A.1, A.2, A.4, 2.1, and A.6. Luckily several approximations are available that enable the solving to be done in steps, instead of solving the whole coupled system. The most general one is discussed next.

The approximation uses the fact that in the case of thermodynamically ideal systems at equilibrium, the complete left hand side of equation A.4 can be neglected. Consider two walls separated from each other with a certain distance, which both have an EDL. When the EDLs do not overlap, e.g. when the distance between the walls is very large, a bulk concentration of the ions of species i, $c_{i,\infty}$ can be defined. Together with the fact that the electric potential far away from both of the surfaces is zero, the solution is given by the *Boltzmann distribution:* $c_i = c_{i,\infty} exp\left(-z_i e \Psi / k_b T\right)$. Substitution of equation A.2 into equation A.1, and using the Boltzmann distribution for the concentration results in:

$$\nabla^2 \Psi = -\frac{1}{\epsilon \epsilon_0} \sum_i c_{i,\infty} exp\left(-\frac{z_i e \Psi}{k_b T}\right) z_i e \tag{A.7}$$

which is the *Poisson-Boltzmann equation.* For a symmetric electrolyte further simplifications are possible, because $z_+ = -z_- = z$, and $c_{+,\infty} = c_{-,\infty} = c_\infty$. At this point it is also convenient to introduce the reduced potential Ψ^* defined by $\Psi^* = ze\Psi / k_b T$. The resulting equation is:

$$\nabla^2 \Psi^* = \kappa^2 sinh\left(\Psi^*\right) \tag{A.8}$$

where:

$$\frac{1}{\kappa} = \sqrt{\frac{\epsilon \epsilon_0 k_b T}{2 z^2 e^2 c_\infty}} \tag{A.9}$$

is the *Debye length* which is the characteristic length of the EDL introduced before. Equation A.8 is non-linear, however it can be made linear when the *Debye-Huckel approximation* is invoked: This approximation uses a first-order Taylor series expansion of *sinh* around $\Psi^* = 0$, resulting in

$$\nabla^2 \Psi = \kappa^2 \Psi \tag{A.10}$$

A.2 Thin Electric Double Layer Approximation

The next approximation is especially useful when simulating microfluidics. The fact that typical Debye lengths, κ^{-1}, are of the order of several tens of nanometres, whereas the typical dimensions of microfluidic devices, e.g. h, are much larger. For these devices $\kappa h \gg 1$ and the large difference between scales can be used to divide the problem into an inner and outer solution. The net charge density distribution varies exponentially from a certain value at the wall to zero in the liquid bulk, and the length scale of this decay is given by the Debye

length. The electric field inside the EDL is mainly governed by wall charges and is in most cases larger than the externally applied field. Outside the EDL, there is no net charge and therefore only the external potential ϕ contributes to the electric field. Together they form the total potential by linear superposition, i.e. $\Psi = \psi + \phi$. Therefore equation A.10 can be divided into

$$\nabla^2\psi = \kappa^2\psi \tag{A.11}$$

$$\nabla^2\phi = 0 \tag{A.12}$$

which are called the inner solution and outer solution respectively. By substituting the net charge density from equation A.1 into the momentum equation, equation A.6, using equation A.11 to rewrite the equation, and using $\vec{E} = -\nabla\phi$ for the externally electric field, the momentum equation can be rewritten to

$$\rho_f \left(\frac{\partial \vec{u}}{\partial t} + (\vec{u} \cdot \nabla)\, \vec{u} \right) = -\nabla p + \mu\nabla^2\vec{u} + \epsilon\epsilon_0\kappa^2\psi\nabla\phi \tag{A.13}$$

This now means that there are the two separate Poisson equations to obtain the (two) potentials, and decoupled from these, the momentum equation can be used to compute the electroosmotic flow.

The approximation above can be extended if only the bulk effect of the presence of the EDL is taken into account. In that case the velocity at the inner-outer plane is only important. This velocity can be obtained through equation A.6, where the net charge density can substituted from equation A.1. Ψ is taken as the inner potential ψ and the electric field $\vec{E}$ is calculated from the external potential ϕ, which is perpendicular to the electric field inside the EDL. Furthermore it can be assumed that the flow is steady, no pressure gradient is applied, and the advection terms can be neglected. This problem can now be solved by using the fact that there is no net transfer of charge from the EDL to the far field, and the velocity gradients are zero at the inner-outer plane. At the wall there is a no-slip boundary condition. Also the fact that the potential difference across the Debye layer from the shear surface between the charged surface and the electrolyte to the far field is specified by the zeta potential, ζ_w, is used. By using these conditions, the velocity at the inner-outer plane is given by the Helmholtz-Smoluchowski relation

$$\vec{u}_{EOF} = -\frac{\epsilon\epsilon_0\zeta_w}{\mu}\vec{E} = \alpha_{os}\vec{E} \tag{A.14}$$

where $\alpha_{os} = -\epsilon\epsilon_0\zeta_w/\mu$ is a certain number that depends only on the surface and electrolyte properties. The Helmholtz-Smoluchowski velocity is in turn used as a boundary condition for the outer solution and can be seen from a bulk point of view as a slip velocity specified at the wall. However, under certain conditions [42] this can be skipped altogether. In that case the local velocity in the bulk is entirely given by the Helmholtz-Smoluchowski relation, eqn. A.14 and the flow pathlines exactly follow the electric field lines.

A.3 Thick Electric Double Layer Approximation

In order to derive the aforementioned approximations, the fact is used that the Debye length is much smaller than any other dimension of the device. In some nanofluidic devices, this approach can not be used anymore because in these devices it is possible to have at least one dimension, e.g. the height of the device, which is of the order of the Debye length ($\kappa h \approx 1$). In this case the electric double layers overlap and as a result interact with each other. Verwey and Overbeek [201] were the first to develop a theoretical model for the potential distribution in this case. This model uses the Boltzmann distribution and therefore using this model for overlapping EDLs is questionable. Other researchers [39, 154] proposed more elaborate models for overlapping EDLs and it was shown [154] that especially in the case when the Debye length is comparable to one of the dimensions of the device, substantial deviations from the theory of Verwey and Overbeek and their model occur. However, consider the situation where the Debye length is much greater than the typical dimension of the device. When the results of others [39, 154, 212] are compared, it can be concluded that more elaborate models all converge to that of the model of Verwey and Overbeek in the case of $\kappa h \ll 1$. Furthermore, when κh is of order 0.2 and lower the linear Poisson-Boltzmann equation can even be used [212]. Therefore, next consider a nanochannel with $\kappa h \ll 1$,

where the typical length (x-direction) and typical width (y-direction) are much larger than the height h (z-direction). When equation A.10 is solved using the boundary conditions $\Psi|_{z=h/2} = \zeta_w$, and $\frac{d\Psi}{dz}|_{z=0} = 0$, the resulting potential distribution is

$$\Psi(z) = \zeta_w \frac{\cosh(\eta\kappa h)}{\cosh\left(\frac{1}{2}\kappa h\right)} \tag{A.15}$$

where $\eta = z/h$ and has values between $-\frac{1}{2}$ and $+\frac{1}{2}$. Directly from equation A.15 it can be seen that when $\kappa h \ll 1$, the potential $\Psi \approx \zeta_w$, because $\cosh \kappa h \approx 1$ in this case. Furthermore, due to the dimensions of the nanochannel, the role of the EDL is much greater in the z-direction than in any other direction. On the other hand, when the height of the nanochannel is everywhere the same, the z-component of the externally applied electric field is everywhere zero, and now only a function of x and y. Assume the flow is steady, no pressure gradient is applied, and the advection terms can safely be neglected. After the net charge density of equation A.1 is substituted into the momentum equation A.6, and equation A.10 is used to rewrite the equation where the ζ potential is used for the potential, the momentum equation is

$$\mu\nabla^2\vec{u} = \epsilon\epsilon_0\kappa^2\zeta_w\vec{E} \tag{A.16}$$

Because of the dimensions of the device, $\partial^2\vec{u}/\partial z^2 \gg \partial^2\vec{u}/\partial x^2$, $\partial^2\vec{u}/\partial y^2$, the momentum equation can be approximated with

$$\frac{\partial^2\vec{u}}{\partial z^2} \approx \frac{\epsilon\epsilon_0\kappa^2\zeta_w}{\mu}\vec{E} \tag{A.17}$$

with boundary conditions $\vec{u}|_{z=h/2} = 0$ and $\partial\vec{u}/\partial z|_{z=0} = 0$, which specify no-slip at the wall and symmetry with respect to the channel mid-plane, the solution for $\vec{u}$ is

$$\vec{u}_{EOF} = \frac{\epsilon\epsilon_0\zeta_w}{\mu}\vec{E}\,(\kappa h)^2\,\frac{1}{2}\left(\eta^2 - \frac{1}{4}\right) \tag{A.18}$$

which is similar to a Poiseuille velocity profile between two infinite plates separated by a distance h. The above approach is similar to the Hele-Shaw approximation [86]. The velocity field at a certain constant z (or averaged with respect to z) is a function in the (x,y) plane only, and with respect to the (x,y) plane, the flow resembles a potential flow. In the classic Hele-Shaw approximation this means that although viscosity is very important, the vertically averaged flow can be regarded as effectively inviscid in the (x,y) plane. In the case of the nanochannel this means that although the EDL is very important, the presence of the EDL can be neglected in the (x,y) plane. Compared to equation A.14, the velocity specified by equation A.18 is just the Helmholtz-Smoluchowski relation multiplied with a certain factor dependent on the value of κh and the z-coordinate. In order to remove the z-coordinate from the equation, the average velocity with respect to z can be used

$$\vec{u}_{avg} = \int_{-1/2}^{+1/2} \vec{u}\,d\eta = -\frac{\epsilon\epsilon_0\zeta_w}{\mu}\vec{E}\frac{(\kappa h)^2}{12} \tag{A.19}$$

which is the electroosmotic velocity an experimental observer is most likely to see.

Appendix B

General $1/r^n$ Expressions for Particle Mesh Methods

The force splitting in the PPPM method is based on the following trivial identity:

$$\frac{1}{r^n} = \frac{f(r)}{r^n} + \frac{1-f(r)}{r^n} \tag{B.1}$$

where the first part is the exponentially decaying short-range term and the second part is the slow varying long-range term. These requirements leave many choices for the function $f(r)$. Essmann et al. [58] proposed such a function. Their idea is based on the (Euler) Gamma function and a simple identity as follows:

$$\Gamma(z) = \int_0^\infty t^{z-1}\exp(-t)\,dt = \lambda^z \int_0^\infty t^{z-1}\exp(-\lambda t)\,dt \tag{B.2}$$

The next step is to split this intergral at $t = \beta$, evaluating the Gamma function at $z = p/2$ and with $\lambda = |\vec{r}|^2 = r^2$. The equation can then be rewritten to:

$$\frac{\Gamma(p/2)}{r^p} = \int_0^{\beta^2} t^{p/2-1}\exp\left(-r^2 t\right)dt + \int_{\beta^2}^\infty t^{p/2-1}\exp\left(-r^2 t\right)dt = S_I + S_{II} \tag{B.3}$$

The integral term S_{II} can be rewritten with a simple substitution $r^2 t = s^2$ which means that $dt = {}^{2s}/_{r^2}ds$ and the lower limit of the integral change to βr:

$$S_{II} = \int_{\beta r}^\infty \left(\frac{s}{r}\right)^{p-2}\exp\left(-s^2\right)\frac{2s}{r^2}ds = \frac{2}{r^p}\int_{\beta r}^\infty s^{p-1}\exp\left(-s^2\right)ds \tag{B.4}$$

Later it will be shown that this is the short range part. The other intergral term, S_I, is a little bit trickier and must be adapted for Fourier space. The following term can be used to do so:

$$\exp\left(-a^2w^2\right) = \frac{\sqrt{\pi}}{a}\int_0^\infty \exp\left(-\frac{\pi^2u^2}{a^2}\right)\exp(-2\pi\imath wu)\,du \tag{B.5}$$

which is the Fourier integral expression for the Gaussian. The next thing is to use the fact that $r^2 = x^2 + y^2 + z^2$ and apply the Gaussian term in all three dimensions with $a^2 = t$. The result is:

$$S_I = \int_0^{\beta^2} t^{p/2-1}\exp\left(-r^2t\right)dt = \pi^{3/2}\int_0^{\beta^2} t^{p/2-5/2}\int_{\mathbb{R}^3}\exp\left(-\frac{\pi^2u^2}{t}\right)\exp(-2\pi\imath u\cdot r)\,d^3\vec{u}dt \tag{B.6}$$

After changing the order of integration this becomes:

$$\int_{\mathbb{R}^3}\left[\pi^{3/2}\int_0^{\beta^2} t^{p/2-5/2}\exp\left(-\frac{\pi^2u^2}{t}\right)dt\right]\exp(-2\pi\imath\vec{u}\cdot\vec{r})\,d^3\vec{u} \tag{B.7}$$

The term between square brackets needs to evaluated first. Substituting of t with s, where $\pi^2 u^2 = ts^2$ and $dt = -2\pi^2 u^2/s^3 ds$ where the lower and upper limit of the integral changes to ∞ and $\pi u/\beta$ are swapped because of the minus sign. The result is:

$$\pi^{3/2} \int_0^{\beta^2} t^{p/2-5/2} \exp\left(-\frac{\pi^2 u^2}{t}\right) dt = 2\pi^{3/2} \left(\frac{\pi \vec{u}}{1}\right)^{p-3} \int_{\frac{\pi u}{\beta}}^{\infty} s^{2-p} \exp\left(-s^2\right) ds \tag{B.8}$$

which for reasons revealed later can also be written as:

$$2\pi^{3/2} \beta^{(p-3)} \left(\frac{\pi \vec{u}}{\beta}\right)^{p-3} \int_{\frac{\pi u}{\beta}}^{\infty} s^{2-p} \exp\left(-s^2\right) ds \tag{B.9}$$

Now all the terms can be added and used to rewrite equation B.3 into:

$$\frac{1}{r^p} = \pi^{3/2} \beta^{(p-3)} \int_{\mathbb{R}^3} f_p\left(\frac{\pi \vec{u}}{\beta}\right) \exp\left(-2\pi\imath \vec{u} \cdot \vec{r}\right) d^3\vec{u} + \frac{g_p(\beta r)}{r^p} \tag{B.10}$$

where

$$f_p(x) = \frac{2(x)^{p-3}}{\Gamma(p/2)} \int_x^{\infty} s^{2-p} \exp\left(-s^2\right) ds \tag{B.11}$$

and

$$g_p(x) = \frac{2}{\Gamma(p/2)} \int_x^{\infty} s^{p-1} \exp\left(-s^2\right) ds \tag{B.12}$$

This equation gives the long-range part that can be solved with the particle mesh method and the short-range part that can be solved with a particle-particle method for any $1/r^p$ type of interaction. The only requirement is to solve the f_p and g_p functions, which in most cases can be done analytically, and compute the Fourier transformed long-range part of the interparticle force, $\tilde{R}(k)$.

Bibliography

[1] J.L.F. Abascal, E. Sanz, and C. Vega. Triple points and coexistence properties of the dense phases of water calculated using computer simulation. *Physical Chemistry Chemical Physics*, 11:556–562, 2009.

[2] J.L.F. Abascal and C. Vega. A general purpose model for the condensed phases of water: TIP4P/2005. *J. Chem. Phys.*, 123:234505, 2005.

[3] F.F. Abraham, J.Q. Broughton, N. Bernstein, and E. Kaxiras. Spanning the continuum to quantum length scales in a dynamic simulation of brittle fracture. *Europhys. Lett.*, 44(6):783–787, 1998.

[4] P. Ahlrichs and B. Dünweg. Lattice-boltzmann simulation of polymer-solvent systems. *Int. J. Mod. Phys. C*, 9(8):1429–1438, 1998.

[5] P. Ahlrichs and B. Dünweg. Simulation of a single polymer chain in solution by combining lattice boltzmann and molecular dynamics. *J. Chem. Phys.*, 111(17):8225–8239, 1999.

[6] O. Aktas and N.R. Aluru. A combined continuum/DSMC technique for multiscale analysis of microfluidic filters. *J. Comp. Phys.*, 178:342–372, 2002.

[7] J. Alejandre and G.A. Chapela. The surface tension of TIP4P/2005 water model using the ewald sums for the dispersion interactions. *J. Phys. Chem.*, 132:014701, 2010.

[8] M.P. Allen and D.J. Tildesley. *Computer Simulation of Liquids*. Oxford Science, London, 1990.

[9] H.C. Andersen. RATTLE: A velocity version of the SHAKE algorithm for molecular dynamics calculations. *J. Comp. Phys.*, 52:24–34, 1983.

[10] A. Arnold and C. Holm. MMM1D: A method for calculating electrostatic interactions in one-dimensional periodic geometries. *J. Chem. Phys.*, 123(14):144103, 2005.

[11] A. Arnold, J. de Joannis, and C. Holm. Electrostatics in periodic slab geometries I. *J. Chem. Phys.*, 117:2496–2502, 2002.

[12] I. Babuška. Error-bounds for finite element method. *Numer. Math.*, 16:322–333, 1970/1971.

[13] O. Bakajin, T.A. Duke, J. Tegenfeldt, C.F. Chou, S.S. Chan, R.H. Austin, and E.C. Cox. Separation of 100-kilobase DNA molecules in 10 seconds. *Anal. Chem.*, 73:6053–6056, 2001.

[14] S. Balasubramanian, C.J. Mundy, and M.L. Klein. Shear viscosity of polar fluids: Molecular dynamics calculations of water. *J. Chem. Phys.*, 105:11190–11195, 1996.

[15] F. Baldessari and J.G. Santiago. Electrophoresis in nanochannels: brief review and speculation. *Journal of Nanobiotechnology*, 4:12–17, 2006.

[16] H.J.C. Berendsen, J.R. Grigera, and T.P. Straatsma. The missing term in effective pair potentials. *J. Phys. Chem.*, 91:6269–6271, 1987.

[17] H.J.C. Berendsen, J.P.M. Postma, W.F. van Gunsteren, and J. Hermans. *Intermolecular Forces*, page 331. Reidel, Dordrecht, 1981.

[18] H.J.C. Berensden, J.P.M. Postma, W.F.V. Gunsteren, A.D. Nola, and J.R. Haak. Molecular dynamics with coupling to an external bath. *J. Chem. Phys.*, 81(8):3684–3690, 1984.

[19] J.D. Bernal and R.H. Fowler. A theory of water and ionic solution, with particular

reference to hydrogen and hydroxyl ions. *J. Chem. Phys.*, 1(8):515–548, 1933.

[20] R.B. Bird, W.E. Stewart, and E.N. Lightfoot. *Transport Phenomena*. Wiley, 2nd edition edition, 2002.

[21] I. Bitsanis, J. J. Magda, M. Tirell, and H. T. Davis. Molecular dynamics of flow in micropores. *J. Chem. Phys.*, 87(3):1733–1750, 1987.

[22] I. Bitsanis, S.A. Somers, H.T. Davis, and M. Tirell. Microscopic dynamics of flow in molecular narrow pores. *J. Chem. Phys.*, 93(5):3427–3431, 1990.

[23] P. Bordat and F. Müller-Plathe. The shear viscosity of molecular fluids: A calculation by reverse nonequilibrium molecular dynamics. *J. Chem. Phys.*, 116:3362–3369, 2002.

[24] P. Bordat, D. Reith, and F. Müller-Plathe. The influence of interaction details on the thermal diffusion in binary lennard-jones liquids. *J. Chem. Phys.*, 115(19):8978–8982, 2001.

[25] M. Born and R. Oppenheimer. Zurquantentheorie der moleküle. *Annalen der Physik*, 84:457–484, 1927.

[26] J. Bourgat, P. Le Tallec, and M. Tidriri. Coupling boltzmann and navier-stokes equations by friction. *J. Comp. Phys.*, 127(2):227–245, 1996.

[27] D. Branton et al. The potential and challenges of nanopore sequencing. *Nature Biotechnology*, 26:1146–1153, 2008.

[28] F. Brezzi. On the existence, uniqueness and approximation of saddlepoint problems arising from lagrangian multipliers. *RAIRO Anal. Numer.*, 8(R- 2):129–151, 1974.

[29] F. Brochard and P.G. de Gennes. Dynamics of confined polymer chains. *J. Chem. Phys.*, 67(1):52–56, 1977.

[30] A. Brodka, J. Koloczek, A. Burian, J.C. Dore, A.C. Hannon, and A. Fonseca. *Journal of Molecular Structure*, 792-793:78–81, 2006.

[31] M. Cabodi, S.W.P. Turner, and H.G. Craighead. Entropic recoil separation of long DNA molecules. *Anal. Chem.*, 74:5169–5174, 2002.

[32] L. Cannavacciuolo, R.G. Winkler, and G. Gompper. Mesoscale simulations of polymer dynamics in microchannel flows. *Europhys. Lett.*, 83(3):34007, 2008.

[33] R. Car and M. Parrinello. Unified approach for molecular dynamics and density-functional theory. *Phys. Rev. Lett.*, 55(22):2471–2474, 1985.

[34] D.Y. Chan and R.G. Horn. The drainage of thin liquid films between solid surfaces. *J. Chem. Phys.*, 83(10):5311–5324, 1985.

[35] T. Chen, B. Smit, and A.T. Bell. Are pressure fluctuation-based equilibrium methods really worse than nonequilibrium methods for calculating viscosities? *J. Chem. Phys.*, 131:246101, 2009.

[36] Y.S. Choi, S.J. Kim, and M.U. Kim. Molecular dynamics of unstable motions and capillary instability in liquid nanojets. *Phys. Rev. E*, 73:016309, 2006.

[37] C.F. Chou, O. Bakajin, S.W. Turner, T.A.J. Duke, S.S. Chan, E.C. Cox, H.G. Craighead, and R.H. Austin. Sorting by diffusion: an asymmetric obstacle course for continuous molecular separation. *PNAS*, 96:13762–13765, 1999.

[38] C.F. Chou, J.O. Tegenfeldt, O. Bakajin, S.Y. Chan, E.C. Cox, N. Darnton, T. Duke, and R.H. Austin. Electrodeless dielectrophoresis of single and double stranded DNA. *Biophysical Journal*, 83:2170–2179, 2002.

[39] A.T. Conlisk and J. McFerran. Mass transfer and flow in electrically charged micro-and nanochannels. *Anal. Chem.*, 74(9):2139–2150, 2002.

[40] S.T. O Connell and P.A. Thompson. Molecular dynamics-continuum hybrid computations: A tool for studying complex fluid flows. *Phys. Rev. E*, 52:5792–5795, 1995.

[41] J.D. Cross, E.A. Strychalski, and H.G. Craighead. Size-dependent DNA mobility in nanochannels. *Journal of Applied Physics*, 102:024701, 2007.

[42] E.B. Cummings, S.K. Griffiths, R.H. Nilson, and P.H. Paul. Conditions for similitude between the fluid velocity and electric field in electroosmotic flow. *Anal. Chem.*, 72(11):2526–2532, 2000.

[43] E.B. Cummings and A.K. Singh. Dielectrophoresis in microchips containing arrays

of insulating posts: theoretical and experimental results. *Anal. Chem.*, 75:4724–4731, 2003.

[44] E.H. Cuthill and J. McKee. Reducing bandwidth of sparse symmetric matrices. *Proc. ACM 24th National Conf., New York*, pages 157–172, 1969.

[45] C-F Dai and R-Y Chang. Molecular dynamics simulation of thread break-up and formation of droplets in nanoejection system. *Molecular Simulation*, 35(4):334–341, 2009.

[46] S. Daoudi and F. Brochard. Flows of flexible polymer solutions in pores. *Macromolecules*, 11(4):751–758, 1978.

[47] T. Darden, D. York, and L. Pedersen. Particle mesh ewald: An N Log(N) method for ewald sums in large systems. *J. Chem. Phys.*, 98(12):10089–10092, 1993.

[48] P.G. de Gennes. *Scaling concepts in polymer physics*. Cornell University Press, Ithaca, 1979.

[49] R. Delgado-Buscalioni and P.V. Coveney. USHER: An algorithm for particle insertion in dense fluids. *J. Chem. Phys.*, 119(2):978–987, 2003.

[50] M. Deserno and C. Holm. How to mesh up ewald sums. I. a theoretical and numerical comparison of various particle mesh routines. *J. Chem. Phys.*, 109(18):7678–7693, 1998.

[51] M. Deserno and C. Holm. How to mesh up ewald sums. II. an accurate error estimate for the p3m algorithm. *J. Chem. Phys.*, 109(18):7694–7701, 1998.

[52] J. Donea and A. Huerta. *Finite Element Methods for Flow Problems*. Wiley, 2003.

[53] W. E and B. Engquist. The heterogeneous multi-scale methods. *Comm. Math. Sci.*, 1(1):87–133, 2003.

[54] W. E, B. Engquist, X. Li, W. Ren, and E. Vanden-Eijnden. Heterogeneous multiscale methods: A review. *Commun. Comput. Phys.*, 2(367), 2007.

[55] J. Eggers. Dynamics of liquid nanojets. *Phys. Rev. Lett.*, 89(8):084502, 2002.

[56] J.C.T. Eijkel and A. van den Berg. Nanotechnology for membranes, filters and sieves. *Lab on a Chip*, 6:19–23, 2006.

[57] A. Einstein. *Investigations on the Theory of the Brownian Movement*. Dover, New York, 1956.

[58] U. Essmann, L. Perera, M.L. Berkowitz, T. Darden, H. Lee, and L.G. Pedersen. A smooth particle mesh ewald method. *J. Chem. Phys.*, 103:8577–8593, 1995.

[59] P. Ewald. Die berechnung optischer und elektrostatischer gitterpotentiale. *Ann. Phys.*, 369:253–287, 1921.

[60] G. De Fabritiis, R. Delgado-Buscalioni, and P.V. Coveney. Energy controlled insertion of polar molecules in dense fluids. *J. Chem. Phys.*, 121(24):12139–12142, 2004.

[61] E.G. Flekkoy, G. Wagner, and J. Feder. Hybrid model for combined particle and continuum dynamics. *Europhys. Lett.*, 52:271–276, 2000.

[62] M. Foquet, J. Korlach, W.R. Zipfel, W.W. Webb, and H.G. Craighead. Focal volume confinement by submicrometer-sized fluidic channels. *Anal. Chem.*, 76:1618–1626, 2004.

[63] D. Frenkel and B. Smit. *Understanding Molecular Simulations: From Algorithms to Applications*. Academic Press, 2nd edition edition, 2002.

[64] J.B. Freund. Electro-osmosis in a nanometer-scale channel studied by atomistic simulation. *J. Chem. Phys.*, 116(5):2194–2200, 2002.

[65] U. Frisch, D. d´Humiéres, B. Hasslacher, P. Lallemand, Y. Pomeau, and J.P. Rivet. Lattice gas hydrodynamics in two and three dimensions. *Complex Systems*, 1:649–707, 1987.

[66] J. Fu, R.B. Schoch, A.L. Stevens, S.R. Tannenbaum, and J. Han. A patterned anisotropic nanofluidic sieving structure for continuous-flow separation of DNA and proteins. *Nature Nanotechnology*, 2:121–128, 2007.

[67] J. Fu, J. Yoo, and J. Han. Molecular sieving in periodic free-energy landscapes created by patterned nanofilter arrays. *Physical Review Letters*, 97:018103, 2006.

[68] M. Fyta, S. Melchionna, S. Succi, and E. Kaxiras. Hydrodynamic correlations in the translocation of a biopolymer through a nanopore: Theory and multiscale simulations. *Phys. Rev. E*, 78:036704, 2008.

[69] A.L. Garcia, J. Bell, W.Y. Crutchfield, and B.J. Alder. Adaptive mesh and algorithm refinement using direct simulation monte carlo. *J. Comp. Phys.*, 154:134–155, 1999.

[70] M.G. Gauthier and G.W. Slater. Sequence effects on the forced translocation of heteropolymers through a small channel. *J. Chem. Phys.*, 128(17):175103, 2008.

[71] C. Gay, P.G. de Gennes, E. Raphaël, and F. Brochard-Wyart. Injection threshold for a statistically branched polymer inside a nanopore. *Macromolecules*, 29(26):8379–8382, 1996.

[72] A. Ghosh, R. Paredes, and S. Luding. Poiseuille flow in a nanochannel - use of different thermostats. In *CD Proceedings, PARTEC 2007, Nuremberg, Germany*, pages 1–4, 2007.

[73] M.A. Gonzàlez and J.L.F. Abascal. The shear viscosity of rigid water models. *J. Chem. Phys.*, 132:096101, 2010.

[74] A. Gopinathan and Y.W. Kim. Polymer translocation in crowded environments. *Phys. Rev. Lett.*, 99:228106, 2007.

[75] E.M. Gosling, I.R. McDonald, and K. Singer. *Mol. Phys.*, 26:1475, 1973.

[76] M.S. Green. Markoff random processes and the statistical mechanics of time-dependent phenomena. II. irreversible processes in fluids. *J. Chem. Phys.*, 22:398–413, 1954.

[77] N. Grønbech-Jensen, R. J. Mashl, R.F. Bruinsma, and W.M. Gelbart. Counterion-induced attraction between rigid polyelectrolytes. *Phys. Rev. Lett.*, 78(12):2477–2480, 1997.

[78] S.R. de Groot and Mazur P. *Non-Equilibrium Thermodynamics*. New York: Dover, 1984.

[79] A. Grzybowski and A. Bródka. Electrostatic interactions in molecular dynamics simulation of a three-dimensional system with periodicity in one direction. *Mol. Phys.*, 100(5):635–639, 2002.

[80] G.J. Guo and Y.G. Zhang. Equilibrium molecular dynamics calculation of the bulk viscosity of liquid water. *Mol. Phys.*, 99(4):283–289, 2001.

[81] N.G. Hadjiconstantinou. Hybrid atomistic-continuum formulations and the moving contact-line problem. *J. Comp. Phys.*, 154:245–265, 1999.

[82] N.G. Hadjiconstantinou and A.T. Patera. Heterogeneous atomistic-continuum representations for dense fluid systems. *International Journal of Modern Physics C*, 8:967–976, 1997.

[83] J. Han and H.G. Craighead. Separation of long DNA molecules in a microfabricated entropic trap array. *Science*, 288:1026–1029, 2000.

[84] R. Hartkamp and S. Luding. Anisotropic lennard-jones fluid in a nanochannel. In *World Congress Particle Technology 6, Nuremberg, 4 pages, CD-Proceedings, March 2009*, 2009.

[85] U. Heinbuch and J. Fischer. Liquid flow in pores: Slip, no-slip, or multilayer sticking. *Phys. Rev. A*, 40(2):1144–1146, 1989.

[86] H.J.S. Hele-Shaw. The flow of water. *Nature*, 58:34–36, 1898.

[87] B. Hess. Determining the shear viscosity of model liquids from molecular dynamics simulations. *J. Chem. Phys.*, 116(1):209–217, 2002.

[88] D.M. Heyes, M. Barber, and J.H.R. Clarke. Molecular dynamics computer simulation of surface properties of crystalline potassium chloride. *J. Chem. Soc., Faraday Trans. II*, 73:1485–1496, 1977.

[89] R.W. Hockney and J.W. Eastwood. *Computer Simulation Using Particles*. IOP, 1988.

[90] C.D. Holcomb, P. Clancy, and J.A. Zollweg. A critical study of the simulation of the liquid-vapor interface of a lennard-jones fluid. *Molecular Physics*, 78(2):437–459, 1993.

[91] W.G. Hoover. Canonical dynamics: Equilibrium phase-space distributions. *Phys. Rev. A*, 31(3):1695–1697, 1985.

[92] H.W. Horn, W.C. Swope, J.W. Pitera, J.D. Madura, T.J. Dick, G.L. Hura, and T. Head-Gordon. Development of an improved four-site water model for biomolecular simulations:TIP4P-Ew. *J. Chem. Phys.*, 120:9665–9678, 2004.

[93] L. Huang and D.E. Makarov. Translocation of a knotted polypeptide through a pore. *J. Chem. Phys.*, 129:121107, 2008.

[94] L.R. Huang, E.C. Cox, R.H. Austin, and J.C. Sturm. Continuous particle separation through deterministic lateral displacement. *Science*, 304:987–990, 2004.

[95] L.R. Huang, J.O. Tegenfeldt, J.J. Kraeft, J.C. Sturm, R.H. Austin, and E.C. Cox. A DNA prism for high-speed continuous fractionation of large DNA molecules. *Nature Biotechnology*, 20:1048–1051, 2002.

[96] J.H. Irving and J.G. Kirkwood. The statistical mechanical theory of transport processes, IV. the equations of hydrodynamics. *J. Chem. Phys.*, 18:817–829, 1950.

[97] J.A. Izaguirre, S. Reich, and R.D. Skeel. Longer time steps for molecular dynamics. *J. Chem. Phys.*, 110(19):9853–9864, 1999.

[98] A. Izmitli, D.C. Schwartz, M.D. Graham, and J.J. de Pablo. The effect of hydrodynamic interactions on the dynamics of DNA translocation through pores. *J. Chem. Phys.*, 128:085102, 2008.

[99] J. de Joannis, A. Arnold, and C. Holm. Electrostatics in periodic slab geometries II. *J. Chem. Phys.*, 117:2503–2512, 2002.

[100] T.B. Jones. *Electromechanics of particles*. Cambridge University Press, 1995.

[101] W.L. Jorgensen, J. Chandrasekhar, J.D. Madura, R.W. Impey, and M.L. Klein. Comparison of simple potential functions for simulating liquid water. *J. Chem. Phys.*, 79:926–935, 1983.

[102] W.L. Jorgensen and J.D. Madura. Temperature and size dependence for monte carlo simulations of TIP4P water. *Mol. Phys.*, 56:1381–1392, 1985.

[103] G. Karniadakis, A. Beskok, and N. Aluru. *Microflows and nanoflows: fundamentals and simulation*. Springer, 2005.

[104] J. Kasianowicz, E. Brandin, D. Branton, and D.W. Deamer. Characterization of individual polynucleotide molecules using a membrane channel. *PNAS*, 93(24):13770–13773, 1996.

[105] W. Kohler and S. Wiegand. *Thermal Nonequilibrium Phenomena in Fluid Mixtures*. Lecture Notes in Physics, Vol. 584. Berlin: Springer, 2002.

[106] J. Koplik, J.R. Banavar, and J.F Willemsen. Molecular dynamics of poiseuille flow and moving contact lines. *Phys. Rev. Lett.*, 60(13):1282–1285, 1987.

[107] J. Koplik, J.R. Banavar, and J.F. Willemsen. Molecular dynamics of fluid flow at solid surfaces. *Phys. Fluids A*, 1:781–794, 1989.

[108] E.M. Kotsalis, I. Hanasaki, J.H.Walther, and P. Koumoutsakos. Non-periodic molecular dynamics simulations of coarse grained lipid bilayer in water. *Computers and Mathematics with Applications*, 59:2370–2373, 2010.

[109] E.M. Kotsalis, J.H. Walther, E. Kaxiras, and P. Koumoutsakos. Control algorithm for multiscale flow simulations of water. *Phys. Rev. E*, 79(4):045701(R), 2009.

[110] E.M. Kotsalis, J.H. Walther, and P. Koumoutsakos. Control of density fluctuations in atomistic-continuum simulations of dense liquids. *Phys. Rev. E*, 76:16709, 2007.

[111] M. Krishnan, I. Monch, and P. Schwille. Spontaneous stretching of DNA in a two-dimensional nanoslit. *Nano Letters*, 5:1270–1275, 2007.

[112] R. Kubo. Statistical-mechanical theory of irreversible processes. i. general theory and simple applications to magnetic and conduction problems. *J. Phys. Soc. Japan*, 12:570–586, 1957.

[113] A.J.C. Ladd. Short-time motion of colloidal particles: Numerical simulation via a fluctuating lattice-boltzmann equation. *Phys. Rev. Lett.*, 70:1339–1342, 1993.

[114] A.J.C. Ladd. Numerical simulations of particulate suspensions via a discretized boltzmann equation. part II. numerical results. *J. Fluid Mech.*, 271:311, 1994.

[115] A.J.C. Ladd and R. Verberg. Lattice-boltzmann simulations of particle-fluid suspensions. *J. Stat. Phys.*, 104:1191–1251, 2001.

[116] O.A. Ladyzhenskaya. *The mathematical theory of viscous incompressible flow.* Gordon and Breach Science, New York, 1969.

[117] L.D. Landau and E.M. Lifshitz. *Electrodynamics of Continous Media, Course of Theoretical Physics, Volume 8.* Pergamon Press, 1960.

[118] S.H. Lee, K. Palmo, and S. Krimm. WIGGLE: A new constrained molecular dynamics algorithm in cartesian coordinates. *J. Comp. Phys.*, 210:171–182, 2005.

[119] S.W. De Leeuw, J.W. Perram, and E.R. Smith. Simulation of electrostatic systems in periodic boundary conditions. I. lattice sums and dielectric constants. *Proc. R. Soc. Lond. A*, 373(1752):27–56, 1980.

[120] J. Lekner. Summation of dipolar fields in simulated liquid-vapour interfaces. *Physica A*, 157(2):826–838, 1989.

[121] J. E. Lennard-Jones. On the determination of molecular fields. *Proc. R. Soc. Lond. A*, 106(738):463–477, 1924.

[122] C.L. Luengo, L.J. van Vliet, B. Rieger, and M. van Ginkel. DIPimage: a scientific image processing toolbox for matlab. Delft University of Technology, http://www.qi.tnw.tudelft.nl/DIPlib/, 1999.

[123] J. Lyklema. *Fundamentals of Interface and Colloid Science, Volume II.* Academic Press Limited, 1995.

[124] Q. Ma and J.A. Izaguirre. Targeted langevin stabilization of molecular dynamics. In *Proceedings of the SIAM Conference on Computational Science and Engineering (CSE'03), San Diego, Calif.*, 2003.

[125] M.W. Mahoney and W.L. Jorgensen. A five-site model for liquid water and the reproduction of the density anomaly by rigid, nonpolarizable potential functions. *J. Chem. Phys.*, 112:8910–8922, 2000.

[126] A.P. Markesteijn, G.O.F. Parikesit, Y. Garini, and J. Westerweel. Numerical simulation of electroosmotic flow in complex micro- and nanochannels. In *Proceedings of the Second International Conference on Transport Phenomena in Micro and Nanodevices, 11-15 June 2006, Il Ciocco Hotel and Conference Center, Barga, Italy*, 2006.

[127] A.P. Markesteijn, O.B. Usta, I. Ali, A.C. Balazs, and J.M. Yeomans. Flow injection of polymers into nanopores. *Soft Matter*, 5:4575–4579, 2009.

[128] S. Matysiak, A. Montesi, M. Pasquali, A.B. Kolomeisky, and C. Clementi. Dynamics of polymer translocation through nanopores: Theory meets experiment. *Phys. Rev. Lett.*, 96:118103, 2006.

[129] R.J. Meagher, J.I. Won, L.C. McCormick, S. Nedelcu, M.M. Bertrand, J.L. Bertram, G. Drouin, A.E. Barron, and G.W. Slater. End-labeled free-solution electrophoresis of DNA. *Electrophoresis*, 26:331–350, 2005.

[130] M. Meche, J. Winkelmann, and J. Fischer. Molecular dynamics simulation of the liquid-vapor interface:the lennard-jones fluid. *J. Chem. Phys.*, 107(21):9264–9270, 1997.

[131] S. Miyamoto and P.A. Kollman. SETTLE: An analytical version of the SHAKE and RATTLE algorithm for rigid water models. *Journal of Computational Chemistry*, 13:952–962, 1992.

[132] A. Mohan, A.B. Kolomeisky, and M. Pasquali. Effect of charge distribution on the translocation of an inhomogeneously charged polymer through a nanopore. *J. Chem. Phys.*, 128:125104, 2008.

[133] M. Moseler and U. Landman. Formation, stability & breakup of nanojets. *Science*, 289(5482):1165–1169, 2000.

[134] B. Muth, M.K. Müller, P. Eberhard, and S. Luding. Collision detection and administration methods for many particles with different sizes. In *DEM07 proceedings CD*, pages 1–18, 2007.

[135] H. Nada and J.P.J.M. van der Eerden. An intermolecular potential model for the simulation of ice and water near the melting point: A six-site model of H2O. *J. Chem. Phys.*, 118(16):7401–7413, 2003.

[136] G. Nagayama and P. Cheng. Effects of interface wettability on microscale flow by molecular dynamics simulation. *International Journal of Heat and Mass Transfer*,

47(3):501–513, 2004.

[137] X. Nie, S. Chen, W. E, and M.O. Robbins. A continuum and molecular dynamics hybrid method for micro- and nano-fluid flow. *J. Fluid Mech.*, 500:55–64, 2004.

[138] M.J.P. Nijmeijer, A.F. Bakker, C. Bruin, and J.H. Sikkenk. A molecular dynamics simulation of the lennard-jones liquid-vapor interface. *J. Chem. Phys.*, 89(6):3789–3792, 1988.

[139] S. Nosé. A unified formulation of the constant temperature molecular-dynamics methods. *J. Chem. Phys.*, 81:511–519, 1984.

[140] A.S. Panwar and S. Kumar. Time scales in polymer electrophoresis through narrow constrictions: A brownian dynamics study. *Macromolecules*, 39:1279–1289, 2006.

[141] G.O.F. Parikesit, V.G. Kutchoukov, A. Bossche, I.T. Young, and Y. Garini. Optical detection of electrokinetically manipulated single molecules in a nanofluidic chip. In *Proc. SPIE*, pages 133–141, 2004.

[142] G.O.F. Parikesit, A.P. Markesteijn, V.G. Kutchoukov, O. Piciu, A. Bossche, J. Westerweel, Y. Garini, and I.T.Young. Electroosmotic flow analysis of a branched U-turn nanofluidic device. *Lab Chip*, 5:1067–1074, 2005.

[143] G.O.F. Parikesit, A.P. Markesteijn, O.M. Piciu, A. Bossche, J. Westerweel, I.T. Young, and Y. Garini. Size-dependent trajectories of DNA macromolecules due to insulative dielectrophoresis in submicrometer-deep fluidic channels. *Biomicrofluidics*, 2:024103, 2008.

[144] D.E. Parry. The electrostatic potential in the surface region of an ionic crystal. *Surf. Sci.*, 49(2):433–440, 1975.

[145] D.E. Parry. Errata: The electrostatic potential in the surface region of an ionic crystal. *Surf. Sci.*, 54:195, 1976.

[146] S. Pennathur, F. Baldessari, M. Kattah, P.J. Utz, and J.G. Santiago. Free-solution oligonucleotide separation in nanoscale channels. *Anal. Chem.*, 79(21):8316–8322, 2007.

[147] S. Pennathur and J.G. Santiago. Electrokinetictransport in nanochannels. 2. experiments. *Anal. Chem.*, 77:6782–6789, 2005.

[148] E. Petersen, B. Li, X. Fang, H. Luo, V. Samuilov, D. Gersappe, J. Sokolov, B. Chu, and M. Rafailovich. DNA migration and separation on surfaces with a microscale dielectrophoretic trap array. *Physical Review Letters*, 98:088102, 2007.

[149] H.L. Pi, J.L. Aragones, C. Vega, E.G. Noya, J.L.F. Abascal, M.A. Gonzalez, and C. McBride. Anomalies in water as obtained from computer simulations of the TIP4P/2005 model: density maxima, and density, isothermal compressibility and heat capacity maxima. *Mol. Phys.*, 107:365–374, 2009.

[150] R. Piazza and A. Parola. Thermophoresis in colloidal suspensions. *J. Phys.: Condens. Matter*, 20:153102, 2008.

[151] H.A. Pohl. *Dielectrophoresis - The behaviour of neutral matter in nonuniform electric fields.* Cambridge University Press, 1978.

[152] P. Polyakov, F. Müller-Plathe, and S. Wiegand. Reverse nonequilibrium molecular dynamics calculation of the soret coefficient in liquid heptane/benzene mixtures. *J. Phys. Chem B*, 112(47):14999–15004, 2008.

[153] M. Porto. Ewald summation of electrostatic interactions of systems with finite extent in two of three dimensions. *Journal of Physics A: Mathematical and General*, 33(35):6211, 2000.

[154] W. Qu and D. Li. A model for overlapped EDL fields. *J. Colloid Interface Sci.*, 224:397–407, 2000.

[155] A. Rahman. Correlations in the motion of atoms in liquid argon. *Phys. Rev. A*, 136:405–411, 1964.

[156] D. Rapaport. *The art of MD simulations.* Cambridge, University Press, 1991.

[157] J. Regtmeier, T.T. Duong, R. Eichhorn, D. Anselmetti, and A. Ros. Dielectrophoretic manipulation of DNA: Separation and polarization. *Anal. Chem.*, 79:3925–3932, 2007.

[158] D. Reith and F. Müller-Plathe. On the nature of thermal diffusion in binary lennard-jones liquids. *J. Chem. Phys.*, 112(5):2436–2443, 2000.

[159] W. Ren. Analytical and numerical study of coupled atomistic-continuum methods for fluids. *J. Comp. Phys.*, 227(2):1353–1371, 2007.

[160] W. Ren and W. E. Heterogeneous multiscale method for the modeling of complex fluids and micro-fluidics. *J. Comp. Phys.*, 204(1):1–26, 2005.

[161] R. Riehn, R.H. Austin, and J.S. Sturm. A nanofluidic railroad switch for DNA. *Nano Letters*, 6:1973–1976, 2006.

[162] B. Rousseau, C. Nieto-Draghi, and J.B. Avalos. The role of molecular interactions in the change of sign of the soret coefficient. *Europhys. Lett.*, 67(6):976–982, 2004.

[163] E. Ruckenstein. Can phoretic motions be treated as interfacial tension gradient driven phenomena? *J. Colloid Interface Sci.*, 83(1):77–81, 1981.

[164] R.E. Rudd and J.Q. Broughton. Concurrent coupling of length scales in solid state systems. *Physica Status Solidi B*, 217(1):251–291, 2000.

[165] J.P. Ryckaert, G. Ciccotti, and Berendsen H.J.C. Numerical integration of the cartesian equations of motion of a system with constraints: Molecular dynamics of n-alkanes. *J. Comp. Phys.*, 23:327–341, 1977.

[166] T. Sakaue. DNA electrophoresis in designed channels. *Eur. Phys. J. E*, 19(4):477–487, 2006.

[167] T. Sakaue, E. Raphaël, P.G. de Gennes, and F. Brochard-Wyart. Flow injection of branched polymers inside nanopores. *Europhys. Lett.*, 72(1):83–88, 2005.

[168] J. Sanz-Serna and M. Calvo. *Numerical Hamiltonian Problems*. Chapman and Hall, London, 1994.

[169] H.A. Schwarz. *Gesammelte mathematische Abhandlungen von H.A. Schwarz. Bd. 1, 2.* Berlin Springer, 1890.

[170] V.B. Shenoy, R. Miller, E.B. Tadmor, D. Rodney, R. Phillips, and M. Ortiz. An adaptive finite element approach to atomic-scale mechanics: the quasicontinuum method. *J. Mech. Phys. Solids*, 47(3):611–642, 1999.

[171] B. Shi and V.K. Dhir. Molecular dynamics simulation of the contact angle of liquids on solid surfaces. *J. Chem. Phys.*, 130:034705, 2009.

[172] B. Shi, S. Sinha, and V.K. Dhir. Molecular dynamics simulation of the density and surface tension of water by particle-particle particle-mesh method. *J. Chem. Phys.*, 124(20), 2006.

[173] H. Shin, M. Oschwald, M.M. Micci, and W. Yoon. Influence of thermodynamic state on nanojet break-up. *Nanotechnology*, 16:2838–2845, 2005.

[174] S.K. Sia and G.M. Whitesides. Microfluidic devices based on poly(dimethylsiloxane) for biological studies. *Electrophoresis*, 24:3563–3576, 2003.

[175] S. Sinha, V.K. Dhir, B. Shi, J.B. Freund, and E. Darve. Surface tension evaluation in lennard-jones fluid system with untruncated potentials, 2003.

[176] V.P. Sokhan, D. Nicholson, and N. Quirke. Fluid flow in nanopores: An examination of hydrodynamic boundary conditions. *J. Chem. Phys.*, 115(8):3878–3887, 2001.

[177] S.A. Somers and H.T. Davis. Microscopic dynamics of fluids confined between smooth and atomically structured solid surfaces. *J. Chem. Phys.*, 96(7):5389–5407, 1992.

[178] Y.M. Song and L.L. Dai. The shear viscosities of common water models by non-equilibrium molecular dynamics simulations. *Molecular Simulation*, 36(7):560, 2010.

[179] E. Spohr. Effect of electrostatic boundary conditions and system size on the interfacial properties of water and aqueous solutions. *J. Chem. Phys.*, 107:6342–6348, 1997.

[180] S.M. Stavis, J.B. Edel, K.T. Samiee, and H.G. Craighead. Single molecule studies of quantum dot conjugates in a submicrometer fluidic channel. *Lab on a Chip*, 5:337–343, 2004.

[181] W.A. Steele. The physical interaction of gases with crystalline solids. *Surf.Sci.*, 36:317–352, 1973.

[182] D. Stein, F.H.J. van der Heyden, W.J.A. Koopmans, and C. Dekker. Pressure-driven transport of confined DNA polymers in fluidic channels. *PNAS*, 103(43):15853–15858, 2006.

[183] C.C. Striemer, T.R. Gaborski, J.L. McGrath, and P.M. Fauchet. Charge- and size-based separation of macromolecules using ultrathin silicon membranes. *Nature*, 445:749–753, 2007.

[184] A.K. Subramanian and C.T. Sun. Continuum interpretation of virial stress in molecular simulation. *International Journal of Solids and Structures*, 45:4340–4346, 2008.

[185] W.C. Swope, H.C. Andersen, P.H. Berens, and K.R. Wilson. A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: Application to small water clusters. *J. Chem. Phys.*, 76:637–649, 1982.

[186] P. Le Tallec and F. Mallinger. Coupling boltzmann and navier-stokes equations by half fluxes. *J. Comp. Phys.*, 136:51–67, 1997.

[187] C. Taylor and P. Hood. A numerical solution of the navier-stokes equations using the finite element method. *Comput. & Fluids*, 1:73–100, 1973.

[188] J.O. Tegenfeldt, C. Prinz, H.Cao, R.L. Huang, R.H. Austin, S.Y. Chou, E.C. Cox, and J.C. Sturm. Micro-and nanofluidics for DNA analysis. *Analytical and Bioanalytical Chemistry*, 378:1678–1692, 2004.

[189] P.A. Thompson and M.O. Robbins. Shear flow near solids: Epitaxial order and flow boundary conditions. *Phys. Rev. A*, 41(12):6830–6837, 1990.

[190] A. Tiwari and J. Abraham. Dissipative particle dynamics simulations of liquid nanojet breakup. *Microfluidics and Nanofluidics*, 4(3):227–235, 2008.

[191] S. Tiwari and A. Klar. Coupling of the boltzmann and euler equations with adaptive domain decomposition procedure. *J. Comp. Phys.*, 144:710–726, 1998.

[192] M.R. Toosi and M.H. Peyrovi. Influence of molecular parameters on the thermodiffusion and thermal conductivity in binary mixtures of diatomic fluids using NEMD. 34:61–74, 2009.

[193] S. Toxvaerd. The structure and thermodynamics of a solid-fluid interface. *J. Chem. Phys.*, 74(3):1998–2005, 1981.

[194] K.P. Travis and K.E. Gubbins. Poiseuille flow of lennard-jones fluids in narrow slit pores. *J. Chem. Phys.*, 112(4):1984–1994, 2000.

[195] K.P. Travis, B.D. Todd, and D.J. Evans. Departure from navier-stokes hydrodynamics in confined liquids. *Phys. Rev. E*, 55(4):4288–4295, 1997.

[196] A. Trokhymchuk and J. Alejandre. Computer simulations of liquid/vapor interface in lennard-jones fluids: Some questions and answers. *J. Chem. Phys.*, 111(18), 1999.

[197] O.B. Usta, A.J.C. Ladd, and J.E. Butler. Lattice-boltzmann simulations of the dynamics of polymer solutions in periodic and confined geometries. *J. Chem. Phys.*, 122:094902, 2005.

[198] J.C. Venter et al. The sequence of the human genome. *Science*, 291:1304–1351, 2001.

[199] L. Verlet. Computer experiments on classical fluids. I thermodynamical properties of lennard-jones molecules. *Phys. Rev.*, 159(98-103), 1967.

[200] L. Verlet. Computer experiments on classical fluids. II equilibrium correlation functions. *Phys. Rev.*, 165(201-214), 1968.

[201] E.J.W. Verwey and J.Th.G. Overbeek. *Theory of the Stability of Lyophobic Colloids*. Elsevier, New York, NY, 1948.

[202] J.L. Viovy. Electrophoresis of DNA and other polyelectrolytes: Physical mechanisms. *Reviews of Modern Physics*, 72:813–872, 2000.

[203] D.C. Wadsworth and D.A. Erwin. One-dimensional hybrid continuum/particle simulation approach for rarefied hypersonic flows. In *AIAA and ASME, 5th Joint Thermophysics and Heat Transfer Conference*, 1990.

[204] J.J. Waterston. On the physics of media that are composed of free and elastic molecules in a state of motion. *Roy. Soc. Proc.*, 5(604), 1846/1893.

[205] E.J.W. Wensink, A.C. Hoffmann, P.J. van Maaren, and D. van der Spoel. Dynamic properties of water/alcohol mixtures studied by computer simulation. *J. Chem. Phys.*, 119(14):7308–7317, 2003.

[206] T. Werder, J.H. Walther, and P. Koumoutsakos. Hybrid atomistic-continuum method

for the simulation of dense fluid flows. *J. Comp. Phys.*, 205:373–390, 2005.

[207] S. Wiegand. Thermal diffusion in liquid mixtures and polymer solutions. *J. Phys.: Condens. Matter*, 16:357–379, 2004.

[208] H.S. Wijesinghe and N.G. Hadjiconstantinou. Discussion of hybrid atomistic-continuum methods for multiscale hydrodynamics. *International Journal for Multiscale Computational Engineering*, 2(2):189–202, 2004.

[209] H.S. Wijesinghe, R.D. Hornung, A.L. Garcia, and N.G. Hadjiconstantinou. Three-dimensional hybrid continuum-atomistic simulations for multiscale hydrodynamics. *J. Fluids Eng.*, 126(768), 2004.

[210] D. Wolf-Gladrow. *Lattice-Gas Cellular Automata and Lattice Boltzmann Models, Vol. 1725, Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2000.

[211] Y.J. Wu, H.L. Tepper, and G.A. Voth. Flexible simple point-charge water model with improved liquid-state properties. *J. Chem. Phys.*, 124(2):024503, 2006.

[212] S. Yao and J.G. Santiago. Porous glass electroosmotic pumps: theory. *J. Colloid Interface Sci.*, 268:133–142, 2003.

[213] I.C. Yeh and M.L. Berkowitz. Ewald summation for systems with slab geometry. *J. Chem. Phys.*, 111:3155–3162, 1999.

[214] S. Yongli, S. Minhua, C. Weidong, M. Congxiao, and L. Fang. The examination of water potentials by simulating viscosity. *Computational Materials Science*, 38(4):737–740, 2007.

[215] I.T. Young, J.J. Gerbrands, and L.J. vanVliet. *Image Processing Fundamentals*, chapter 51 of The Digital Signal Processing Handbook. CRC Press LLC, 1998.

[216] L. Zhang, R. Balasundaram, S. Jiang, and S.H. Gehrke. Nonequilibrium molecular dynamics simulation of confined liquids in contact with the bulk. *J. Chem. Phys.*, 114(15):6869–6877, 2001.

[217] O.C. Zienkiewicz, R.L. Taylor, and J.Z. Zhu. *The Finite Element Method: Its Basis and Fundamentals*. Elsevier, 6th edition edition, 2005.

[218] M. Zwolak and M. Di Ventra. Physical approaches to DNA sequencing and detection. *Rev. Mod. Phys.*, 80(1):141–165, 2008.

List of Publications

Journal Contributions

- G.O.F. Parikesit, A.P. Markesteijn, V.G. Kutchoukov, O.M. Piciu, A. Bossche, J. Westerweel, Y. Garini, and I.T. Young, Electroosmotic flow analysis of a branched U-turn nanofluidic device. *Lab on a Chip*, 5(10):1067, 2005
- V.H.J. Nieborg, A.P. Markesteijn, R.H. Lindken, G.J. Witkamp, H.J.M. Kramer, and J. Westerweel, "Mixing with EOF for a Precipitation Reaction. *Journal of Dispersion Science and Technology*, 29(4):587, 2008
- G.O.F. Parikesit, A.P. Markesteijn, O.M. Piciu, A. Bossche, J. Westerweel, I.T. Young, and Y. Garini, Size-dependent trajectories of DNA macromolecules due to insulative dielectrophoresis in submicrometer-deep fluidic channels. *Biomicrofluidics*, 2(2), 2008
- A. P. Markesteijn, O. B. Usta, I. Ali, A. C. Balazs, and J. M. Yeomans, Flow injection of polymers into nanopores. *Soft Matter* 5(22):4575, 2009
- L.C. Jellema, A.P. Markesteijn, J. Westerweel and E. Verpoorte, Tunable Hydrodynamic Chromatography of Microparticles Localized in Short Microchannels. *Anal. Chem.*, 82(10):4027, 2010
- A.P. Markesteijn, R.M. Hartkamp, S. Luding, and J. Westerweel, A comparison of the value of viscosity for several water models using Poiseuille flow in a nano channel, 2011, Submitted

Conference Contributions

- A.P. Markesteijn, G.O.F. Parikesit, Y. Garini, and J. Westerweel, Numerical simulation of electroosmotic flow in complex micro- and nanochannels. In *Proceedings of the Second International Conference on Transport Phenomena in Micro and Nanodevices, 11-15 June 2006, Il Ciocco Hotel and Conference Center, Barga, Italy*, 2006
- A.P. Markesteijn and J. Westerweel, Connecting Molecular Dynamics and Computational Fluid Dynamics. In *Book of Abstract of the "16th Discrete Simulation of Fluid Dynamics (DSFD), July 23-27, 2007, Banff, Canada*, 2007
- A.P. Markesteijn and J. Westerweel, Connecting Molecular Dynamics and Computational Fluid Dynamics., In *Book of Abstract of the "IUTAM Symposium on Advances in Micro- and Nanofluidics, September 6-8, 2007, Dresden, Germany*, 2007
- A.P. Markesteijn and J. Westerweel, Connecting Molecular Dynamics and Computational Fluid Dynamics, In *Book of Abstracts of the "Dutch MicroNano Conference 2007, Wageningen, The Netherlands*, 2007
- A.P. Markesteijn and J. Westerweel, Connecting Molecular Dynamics and Computational Fluid Dynamics, Flow round nano-sized objects, investigating non-continuum effects, and new boundary conditions for MD. In *Book of Abstracts of the Physics@Veldhoven meeting, Foundation for Fundamental Research on Matter (FOM), Veldhoven, The Netherlands*, 2008
- A.P. Markesteijn and J. Westerweel, Connecting Molecular Dynamics and Computational Fluid Dynamics, Flow round nano-sized objects, investigating non-continuum

effects, and new boundary conditions for MD. In *Book of Abstracts of The Fifth ICMMES, June, 16-20, 2008, Amsterdam, The Netherlands*, 2008

Posters, Presentations, and Other Contributions

- A.P. Markesteijn, G.O.F. Parikesit, Y. Garini, and J. Westerweel, Numerical simulation of electroosmotic flow in complex micro- and nanochannels. In *Book of Abstracts of the Foundation for Fundamental Research on Matter (FOM) meeting 2005, Veldhoven, The Netherlands*, 2005
- G.O.F. Parikesit, A.P. Markesteijn, O.M. Piciu, V.G. Kutchoukov, J. Westerweel, A. Bossche, Y. Garini, and I.T. Young, Numerical simulation of electrokinetic force fields for particle manipulation and sorting in a branched U-turn 2D-like nanofluidic device. In *Proceedings of μTAS (Miniaturized Systems for Chemistry and Life Sciences) Conference, Boston, USA, Vol. 2, pages 841-843*, 2005
- G.O.F. Parikesit, A.P. Markesteijn, O.M. Piciu, V.G. Kutchoukov, J. Westerweel, A. Bossche, Y. Garini, and I.T. Young, The Molecular Pinball Machine: A novel nanofluidic device for single-molecule analysis. In *Book of Abstracts of The 20th International Symposium on Micro-Scale Bioseparations (MSB), Amsterdam, The Netherlands*, 2006
- G.O.F. Parikesit, A.P. Markesteijn, O.M. Piciu, V.G. Kutchoukov, J. Westerweel, A. Bossche, Y. Garini and I.T. Young, Electrokinetics in a branched U-turn nanofluidic device. In *Book of Abstracts of the Physics@Veldhoven meeting, Foundation for Fundamental Research on Matter (FOM), Veldhoven, The Netherlands*, 2007
- G.O.F. Parikesit, A.P. Markesteijn, O.M. Piciu, V.G. Kutchoukov, J. Westerweel, A. Bossche, Y. Garini, and I.T. Young, Sorting of DNA single-molecules in a branched nanofluidic U-turn. In *Book of Abstract of the Physics and Chemistry of Microfluidics Meeting, Gordon Research Conferences (GRC), Waterville Valley, USA*, 2007
- G.O.F. Parikesit, A.P. Markesteijn, J. Westerweel, I.T. Young, and Y. Garini, Free-Flow Dielectrophoresis-A numerical study. In *Proceedings of μTAS (Miniaturized Systems for Chemistry and Life Sciences) Conference, Paris, France, Vol. 2, pages 221-223*, 2007
- G.O.F. Parikesit, A.P. Markesteijn, O.M. Piciu, V.G. Kutchoukov, J. Westerweel, A. Bossche, Y. Garini, and I.T. Young, Free-Flow Dielectrophoresis. In *Book of Abstracts of the Physics@Veldhoven meeting, Foundation for Fundamental Research on Matter (FOM), Veldhoven, The Netherlands*, 2008
- L.C. Jellema, A.P. Markesteijn, J. Westerweel, and E.M.J. Verpoorte, Size Separation of Particles Using Flow-Induced Electrokinetic Trapping. In *Proceedings of μTAS 2008 (Miniaturized Systems for Chemistry and Life Sciences) Conference, San Diego, California, USA, October 12-16*, 2008

Acknowledgments

This PhD thesis forms the conclusion of my research at the Laboratory for Aero- and Hydro-Dynamics at the University of Technology Delft. It has been an interesting period with a steep learning curve. Several people have made it possible for me to achieve this goal.

First of all, none of this research would have been possible, hadn't it been for the funding from MicroNED and the EU INFLUS-project. I would also like to thank my promoters Jerry Westerweel and Stefan Luding for their guidance, support, and patience. I would like to thank Jerry for his ideas and directions, especially at the beginning of the research. I would also like to thank him for the extracurricular opportunities he made possible. I would like to thank Stefan for pointing me in the right direction with regards to the Molecular Dynamics.

I would like to thank my colleagues Gea Parikesit, Sabeth Verpoorte and her department at the University of Groningen, and Remco Hartkamp. From the start of my research, Gea provided me with the practical side of microfluidics and nanofluidics, which helped me with the initial stages of my research. He has a wonderful sense of humour. I would like to thank "Groningen" for their challenging points of views and our interesting discussions. In the last stage of my PhD, Remco proved to be a valuable sounding board.

Further, I would like to thank my colleagues of the Lab, especially my roommates Chiara Tesauro and Peter Vennemann. I was always nice to chat. I would like to thank Chiara for introducing me to the wonderful espresso chocolate sweets that were a good substitute for coffee. I am grateful to Ria and Caroline for chasing after the paperwork, especially in the final stages of my PhD.

Special thanks go to Julia Yeomans and the "Condensed Matter Theory: Soft and Biological Matter" department in Oxford. They welcomed me during my six months there and made me feel at home. I would like to thank Julia, Anna Balazs, and Berk Usta for further familiarising me with the Lattice Boltzmann method and giving me the opportunity to work with the computer modelling of polymers.

Last but not least, I would like to thank my family, especially my wife Johanna and my daughter Melisse. Johanna: Thank you for your patience and feedback. Melisse: Thank you for providing welcome diversions.

About the Author

Anton Pieter Markesteijn was born on 26 January 1980 in Papendrecht, The Netherlands. He attended the Willem de Zwijger in Papendrecht doing a Baccalauréat type A-levels (VWO) in Dutch, English, Mathematics, Physics, Chemistry and Economics, starting in 1992 and graduating in 1998. That same year, he started his academic education at Delft University of Technology in the department of Mechanical Engineering. During this period he also worked for two small engineering companies developing software. In the summer of 2004, he obtained his Master of Science degree in 'Mechanical Engineering - Solid and Fluid Mechanics', specialisations 'Engineering Dynamics' and 'Fluid Mechanics', both with Honours. The project of this M.Sc. thesis was performed at OCE Venlo at the Research & Development department, where he was involved in the droplet formation process, and the fluid and structural modelling of a new type of print head used for ink-jet printing.

In the autumn of 2004, Anton started the research project which led to this thesis performed in the Laboratory of Aero- and Hydrodynamics at Delft University of Technology. This work started under the supervision of Prof. Dr. Ir. Jerry Westerweel and adviser Dr. rer. nat. S. Luding. The goal of this research was to better understand and develop a computational method where molecular dynamics and a continuum method are coupled, which can be used for industrial applications. In April 2008 he was invited to do a six month research project at the University of Oxford, in the Oxford Theoretical Physics Condensed Matter Theory: Soft and Biological Matter division under supervision of Prof. Julia Yeomans and adviser Dr. Berk Usta. The goal of this project was to investigate the flow injection of polymers into nanopores, which led to part of this thesis.

On 20 May 2006 he married Johanna Fusfeld and their lovely daughter Melisse was born on the 20 October 2006. In May 2008 he moved to Ringshall, Suffolk, Great Britain where he is currently living.