1 Introduction

Hydrodynamics describes fluid dynamical properties like mass density, streaming velocity and energy [15]. Understanding and controlling fluids has a long history [50], and despite the prefix, hydrodynamics extends today beyond the study of water, and more generally to all fluids. The definition of a fluid is ambiguous, but we here think of it in an intuitive manner as a material (or substance) that will easily flow when a force is applied to it. Liquids and gasses are two usual examples of a fluid.

The fundamental assumption in hydrodynamics is that the properties vary sufficiently smoothly in both time and space; this is known as the continuum hypothesis, see, for example, Ref. [141]. In this way the properties can be treated mathematically as field variables. The hypothesis is in agreement with our everyday experience: when we are stirring a cup of coffee or riding a bicycle, the fluid flow appears smooth. However, the continuum hypothesis is not strictly true. Think of a small fluid volume (or fluid element), denoted $\mathcal{V}$, embedded in a material at rest. Hydrodynamics will predict that the mass is constant with respect to time; however, due to thermal motion, molecules will enter and leave the fluid element, and the mass of $\mathcal{V}$ will fluctuate. As the element volume increases or if we perform a sufficiently long time average, the mass of $\mathcal{V}$ will converge to that of the hydrodynamic prediction. What defines sufficient large volumes and time averages is not clear, and we return to this example later in this chapter.

Channels and tubes for fluid flows with nanoscale cross section and nanoscale volume can now be fabricated with impressive accuracy [88, 170]. In order to control and utilise such nanoscale fluid volume devices, it is important to develop models that can describe and predict the fluid dynamics. Is the continuum picture, which has been applied with great success to micro- and macro-fluid systems, also applicable on the nanoscale? From the preceding example, it appears that the answer is not trivially 'yes' or 'no', and this question is the underlying theme of the text.

We must settle on a few important definitions. Eikel and van den Berg [61] define *nanofluidics* as the study and application of fluid flows confined in and around nanosized structures. For fluids confined in nanoscale structures (or geometries) a non-negligible fraction of the fluid molecules will interact with the wall atoms. To describe such systems in detail an in-depth knowledge of both the wall–fluid and fluid–fluid interactions are needed [29, 117, 165]. The complexity can be overwhelming and appear almost intractable from a modelling point of view. Fortunately, many of the underlying physical mechanisms relevant in confined fluids are also present in the non-confined case, and one can simplify the problem considerably by studying these systems

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

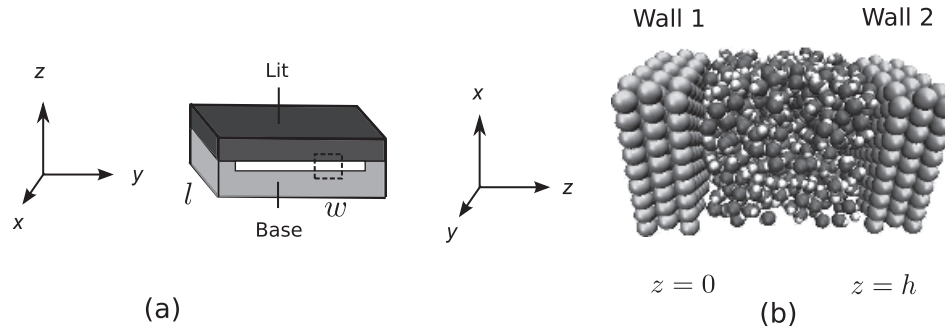

Figure 1.1 (a) Illustration of a slit-pore geometry fabricated by etching a silicon base and placing a lit on the cavity. (b) A snapshot of a computer simulation of water flowing in a region of the slit-pore, highlighted by the punctured box in (a). In the simulations, the walls and fluid are infinite in extent in the (x, y)-plane. The pore height h is 3.4–3.5 nm. From Ref. [102] with permission.

[49, 66, 110, 199]. One can even study some of the mechanisms for non-confined systems in equilibrium [31, 90], that is, in no-flow situations. Therefore, in order to explore nanofluidic systems in detail, we here extend the scope of the book and define *nanoscale hydrodynamics* as

> *the study and application of fluid systems, where the system characteristic length scale is in order of nanometres.*

Whether the system is confined or not, a *nanoscale fluid system* is then a system where the characteristic length scale is in order of nanometres, that is, 1–100 nanometres. Naturally, the length scale needs to be clearly defined for each system. Notice that the related field *nanofluids* is the study of nanosized particles suspended in fluids and is not the focus of the text.

An example of a nanoscale fluid system that we will explore is water confined in a slit-pore, where the pore height, h, is around 10 water diameters or 3.4–3.5 nm. In brief, the slit-pore geometry can be fabricated experimentally from etching a base (typically a silicon wafer) and placing a lit on the formed cavity; see Fig. 1.1 (a). Figure 1.1 (b) is a snapshot from a molecular dynamics, MD, simulation that simulates only a small region of the slit-pore and after the coordinate system has been rotated twice. We have that $w \gg h$ and $l \gg h$, and h is the characteristic length.

We will often use this simple geometry in our study of confined nanoscale fluids. Compared to clever choices of coordinate system, where, for example, the walls are located at dimensionless points $z = \pm 1$, the coordinate system shown in Fig. 1.1 (b) leads to slightly more complicated mathematical expressions. However, the results will depend explicitly on h, the characteristic length scale, and we will stick with this more intuitive choice of coordinate system with the exception of Section 5.6, where we explore molecular fluid flows.

In confinement one can decompose the forces acting in the system into surface forces (e.g., the wall-fluid frictional force) and volume forces (e.g., the gravitational force). The

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

total surface force and volume force must be proportional to the wall–fluid surface area and fluid volume, respectively. For the geometry in Fig. 1.1 we have

$$\frac{\text{Total surface force}}{\text{Total volume force}} \propto \frac{l \times w + h \times l}{w \times l \times h} \approx \frac{1}{h}, \tag{1.1}$$

as $h \ll w$. This is also known as the square-cube law [1] and shows that the surface force becomes dominant for small characteristic lengths. One immediate result of this is that confined nanoscale fluid flows cannot be generated by, say, Earth's gravitational pull, but must be realised through application of more advanced methods such as electro-osmosis.

In nanofluidic laboratory experiments the fluid velocity rarely exceeds 0.1 m/s; see Whitby and Quirke [212], corresponding to a flow rate less than 10^{-15} L/s. If we define the Reynolds number, Re, as

$$\mathrm{Re} = \rho h U / \eta_0 \,, \tag{1.2}$$

where ρ is the mass density, $U = 0.1$ m/s is the characteristic fluid velocity, and η_0 is the shear viscosity, the Reynolds number is in practise below 0.01. The flow is therefore a Stokes flow, or creeping flow, and we can safely neglect the advective inertial forces in the hydrodynamic description of the system.

Molecular dynamics is widely used today to study nanoscale fluid systems. Common for these computer-based studies is that unrealistically large flow velocities are simulated. Usually the velocity is in order of 10–10^2 m/s, nevertheless, due to the small length scales, the Reynolds number is still small, often between 1 and 10, and inertial forces can again be ignored when analysing the simulation data.

Another important point about simulation studies is that despite the very large fluid velocities, the flow speed is usually significantly lower than the corresponding sound speed, c_s. This is commonly quantified through the Mach number,

$$\mathrm{Ma} = U / c_s \,. \tag{1.3}$$

If $\mathrm{Ma} < 0.3$, the fluid compressibility effects can usually be ignored. For water at ambient conditions, c_s is on the order of 10^3 m/s and the corresponding Mach number is below 0.1 in simulations under usual simulation conditions. Nanoscale flow systems, both real and simulated, are thus characterised as being laminar and incompressible, and this simplifies the mathematical analysis of the models considerably, as we will see.

In our exploration of nanoscale hydrodynamics we focus on fluid systems characterised by a system relaxation time that is sufficiently small compared to the time scale at which we perform measurements or simulations. Specifically, if τ_s is the time for the fluid internal stress to relax after shearing/deformation and τ_{obs} is the time we observe the system, we define the Deborah number

$$\mathrm{De} = \tau_s / \tau_{\mathrm{obs}} \,, \tag{1.4}$$

which then must be significantly smaller than 1. For fluids like methane, butane, and water at ambient conditions, a small Deborah number is easily obtained, even in computer simulations. On the other hand, polymer and glass systems feature very large

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

relaxation times having a large Deborah number even if τ_{obs} is the real laboratory observation time.

We here refer to fluids with sufficiently small τ_s compared to τ_{obs} as simple fluids; τ_{obs} is often defined by what can be achieved in computer simulations. By simple systems we mean simple fluids, either unconfined or confined to simple straight channels like slit-pores or nanochannels.

Nanofluidics is believed to play a critical role in many areas of future engineering [61, 132, 176]. The purpose here is not to present the many exciting applications, but to show how nanoscale fluid systems can be modelled and what new insight into fluid and liquid theory this brings. Often what appears to be a new phenomenon specific to the nanoscale is actually omnipresent, but can be ignored on larger length scales. The term *nanoscale fluid phenomenon* is still used when the phenomenon is particularly relevant to nanoscale fluid systems.

1.1 Nanoscale Fluid Phenomena

Before going into detail on how to model nanoscale fluid systems and what we can learn from that, it will be enlightening to see a few examples of some of the phenomena we will explore later. The examples will by no means cover all phenomenology, but they illustrate at least some unique features of these systems and also motivate the topics in the book.

1.1.1 Flows in Nanochannels

Horn and Isrealachvili [114] showed in 1981 that the force acting between two mica surfaces in liquid octamethyl-cyclotetra-siloxane (OMCTS) features oscillations as the distance between the mica surfaces is varied. Specifically, the oscillations have a period matching the diameter of OMCTS, and the amplitude decay as the distance between the surfaces is increased. This seminal result is a fingerprint of a molecular layering near the surface, a layering which is strong in the fluid region close to the surface and decays with distance. This layering was also reported in 1977 by Toxvaerd and Præstgaard [202] who studied confined systems using molecular dynamics simulations. The focus here is on the hydrodynamics, and we therefore ask how the layering affects the fluid flow properties. For simple fluids, as we have defined it above, the effect is surprisingly small, but for non-simple fluids composed of, say, long alkane chains the picture is much more complicated, and the fluid's flow resistance increases significantly as we reach nanoscale confinement [83]. This increase is attributed to adhesion/cohesion effects coming from molecular layering, but also to crystallisation, vitrification, phase transitions, and more; see a summary in Ref. [68]. The increased flow resistance has led to the concept of an effective viscosity, η_{eff}, which in nanoscale systems can be many times larger than the shear viscosity η_0 characterising the flow properties in the macroscopic or non-confined case.

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

The opposite effect due to confinement, namely a flow enhancement, has also been observed. This phenomenon can be quantified through the enhancement coefficient, E, which is defined as the ratio between the experimentally measured volumetric flow rate, Q_{exp}, and the theoretical predicted flow rate, Q_{the},

$$E = Q_{\text{exp}}/Q_{\text{the}}\,. \tag{1.5}$$

Recall, the volumetric flow rate is the fluid volume discharged by the channel per time unit [38]. The predicted volumetric flow rate Q_{the} is often calculated from the Navier–Stokes equation in the given geometry, using η_0, and with specified boundary conditions; thus, we assume that the Navier–Stokes equation is applicable on the nanoscale.

In 2005, Majumder et al. [153] investigated water flow through carbon nanotubes embedded in a membrane and having diameters of around 7 nm. Using zero velocity boundary conditions to calculate Q_{the}, the authors reported a surprisingly large enhancement coefficient, on the order of 10^4. One hypothesis for the mechanism behind the enhancement is that the fluid velocity at the wall–fluid boundary is nonzero. This is referred to as slippage and can be quantified from the slip enhancement coefficient,

$$E^{\text{slip}} = Q_{\text{the}}^{\text{slip}}/Q_{\text{the}}^{\text{noslip}}\,, \tag{1.6}$$

where $Q_{\text{the}}^{\text{slip}}$ is the theoretical prediction for the volumetric flow rate using slip boundary conditions and $Q_{\text{the}}^{\text{noslip}}$ is the prediction using the traditional no-slip boundary conditions. For a steady flow in slit-pore geometries (Fig. 1.1), the slip-enhancement coefficient can be evaluated to

$$E^{\text{slip}} = 1 + 6L_s/h\,, \tag{1.7}$$

where L_s is the slip length. The slip length is the interesting quantity when discussing flow enhancement, and is, for the planar Poiseuille flow, the distance away from the wall–fluid interface to where the linearly extrapolated velocity is zero. The slip length is illustrated in Fig. 1.2 (a).

For highly hydrophobic walls, the slip length is expected to be large; the results for carbon nanotubes from Majumder et al. suggest L_s to be on the order of 10^4 nm. The actual slip length magnitude is still debated. For example, from molecular dynamics simulations Kannam et al. [125] found a slip length of around 100 nm in a carbon nanotube with a diameter of 4 nm; this corresponds to a plug-like flow. From experiments on a Landau–Squire flow, Secchi et al. [188] found the same slip length, but for carbon nanotubes with diameters of around 50 nm. There is general agreement that the slip length decreases as the tube diameter increases and that it converges to that of graphene; experimental, theoretical, and simulation studies find $L_s = 10$–80 nm for graphene–water; see Ref. [125] and references therein. Figure 1.2(b) plots the slip enhancement coefficient for graphene–water as a function of slit-pore height h.

Importantly, fluid slippage also occurs on macroscopic length scales, but from Eq. (1.7) the effect of this on the flow rate is not observed under usual macroscopic

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

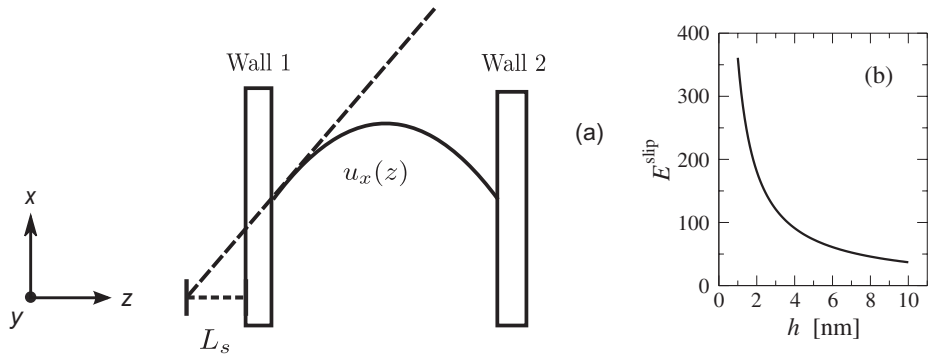

Figure 1.2 (a) Illustration of the slip length, L_s, for the planar Poiseuille flow. The dashed line is the tangent line for the fluid velocity at $z = 0$. The curve illustrates the fluid velocity x-component. (b) Slip enhancement coefficient for a slit-pore where $L_s = 60$ nm (graphene–water).

circumstances, as h is many orders of magnitude larger than the slip length. Even for microscopic length scales, E^{slip} is close to unity.

1.1.2 Capillary Raise

A common way to fill nanochannels and nanopores with a fluid is through capillary filling and, naturally, this method has drawn a lot of attention to the research community. Capillary filling in micro- and macropores is, under usual conditions, described satisfactory by the Lucas–Washburn equation, however, on the nanoscale the slippage phenomenon introduced in the previous section becomes important. The modified Lucas–Washburn equation including this effect reads [57, 119]

$$h_{\text{cap}}^2(t) = \frac{\gamma R \cos(\theta)}{2\eta_0}\left(1 + \frac{4L_s}{R}\right)t, \tag{1.8}$$

where h_{cap} is the capillary height, γ is the surface tension, R is the radius of the tube, and θ is the contact angle between the meniscus and the wall; see Fig. 1.3 (a).

Figure 1.3(b) shows the capillary height h_{cap} as a function of time for two different model fluids investigated using molecular dynamics (MD) simulations. It is worth noting that the typical unit length scale in molecular dynamics is 3–5 Å and the unit time scale is on the order of picoseconds. After a short inlet transient time, the capillary height predicted by the modified Lucas–Washburn equation is confirmed – even quantitatively.

The research continues, as questions remain unresolved. For example, it is known that the contact angle and the slip length are correlated quantities [211], and this calls for a revision of the fundamental theory behind capillary filling on the nanoscale.

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

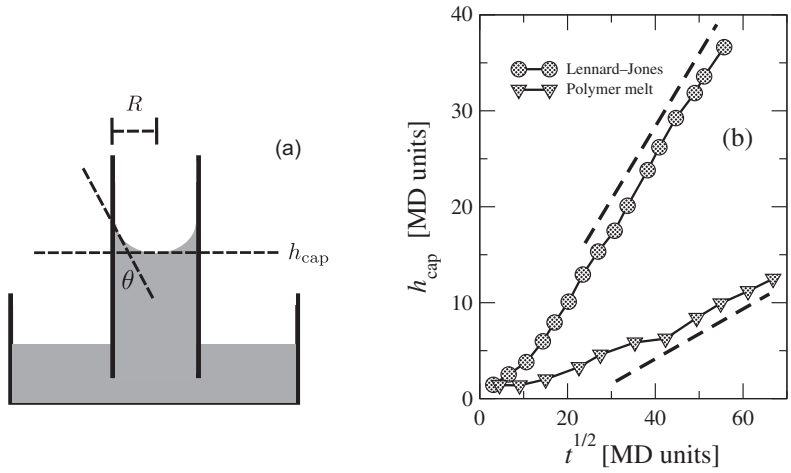

Figure 1.3 (a) Illustration of capillary raise. (b) Capillary height as a function of time in a tube of radius 10 (in MD units); filled circles are data for a Lennard–Jones liquid and triangles data for a model polymer melt. Punctured lines illustrate the theoretical predictions for the slopes. From Dimitrov et al. [57].

1.1.3 Anisotropy in Confined Dielectrics

Application of an external electrical field, $\mathbf{E}^{\mathrm{ext}}$, to a dielectric material gives rise to a polarisation, $\mathbf{P}$. Recall, in the static, homogeneous, isotropic, and linear cases, the polarisation is

$$\mathbf{P} = \varepsilon_0(\varepsilon_r - 1)\mathbf{E}^{\mathrm{ext}} = \varepsilon_0\chi_e\mathbf{E}^{\mathrm{ext}}, \tag{1.9}$$

where ε_r and ε_0 are the relative and vacuum dielectric permittivities, respectively, and $\chi_e = \varepsilon_r - 1$ is the electric susceptibility. In the situation where an external electric field is suddenly switched on at, say, $t = 0$, the system relaxation response can also be studied. The simplest model for this non-static case is the Debye model,

$$\mathbf{P}(t) = \varepsilon_0\chi_e\mathbf{E}^{\mathrm{ext}}(1 - e^{-t/\tau_D}); \tag{1.10}$$

τ_D is the Debye relaxation time. Notice that the polarisation $\mathbf{P}$ converges to $\varepsilon_0\chi_e\mathbf{E}^{\mathrm{ext}}$ in accordance with the static case, Eq. (1.9).

Figure 1.4 shows molecular dynamics results for the polarisation as a function of time for water confined in a slit-pore. As the field is applied parallel to the walls, the system response is bulk-like and follows an exponential relaxation in accordance with the Debye model. On the other hand, the response is significantly changed when the field is applied normal to the walls; here it resembles a small amplitude step response. In both experiments [76] and simulations [149] this anisotropy is found for slit-pore heights up to 100 nm, and we approach the microfluidic length scale.

The reduced polarisation phenomenon indicates that there exist a parallel permittivity $\varepsilon_r^{\parallel}$ and a normal permittivity $\varepsilon_r^{\perp}$ with respect to the wall plane. If the system's dielectric response is isotropic, we have $\varepsilon_r^{\parallel} = \varepsilon_r^{\perp} = \varepsilon_r$.

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

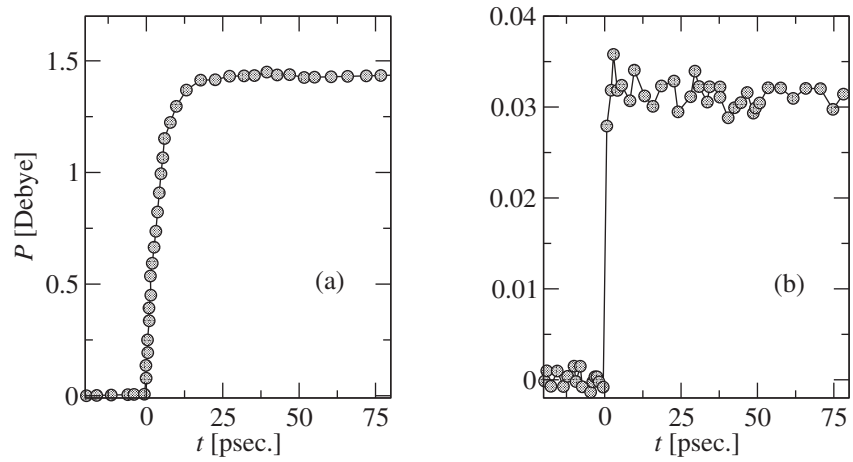

Figure 1.4 Molecular dynamics results for the polarisation as a function of time for water confined in a slit-pore with $h = 11$ nm; see the geometry in Fig. 1.1 (b). (a) The external electric field is applied parallel to the wall at $t = 0$. Same as (a), but where the field is applied normal to the wall. Data are from Ref. [149].

To model the reduced normal permittivity, we can divide the confined fluid into layers with respect to the z-direction. These layers can be considered as capacitors in a serial arrangement. Each capacitor has capacitance C_i, permittivity $\varepsilon_{r,i}^{\perp}$, and separation (or height) h_i. The fundamental idea now is that the capacitance in the wall–fluid interface is low, as the water density is low. The total capacitance is

$$\frac{1}{C} = \sum_i \frac{1}{C_i} \,. \tag{1.11}$$

The capacitance for each layer is $C_i = \varepsilon_{r,i}^{\perp}\varepsilon_0 A/h_i$, where A is the surface area of the capacitor, and we get

$$\varepsilon_r^{\perp} = \frac{h}{\sum_i h_i/\varepsilon_{r,i}^{\perp}} \,. \tag{1.12}$$

Zhang [215] studied the simplest scenario of a fluid next to a wall and included only two fluid layers: one layer (denoted layer 1) just adjacent to the wall, where the height was h_1, and a layer 2 with height h_2, such that $h = h_1 + h_2$. If h_1 is on the order of angstroms, then $\varepsilon_{r,1}^{\perp} \approx 1$, as this interfacial region is almost a vacuum. The second layer, we assume, has permittivity $\varepsilon_{r,2}^{\perp} \approx \varepsilon_r$. Then Eq. (1.12) simplifies to $\varepsilon_r^{\perp} = h/(h_1 + h_2/\varepsilon_r)$. As $h \to h_1$ we have $h_2 \to 0$ and $\varepsilon_r^{\perp} \to 1$. As h increases we reach the regime where $h_2 \gg h_1$ and $h_1 + h_2/\varepsilon_r \approx h_2/\varepsilon_r$, that is, $\varepsilon_r^{\perp} \approx \varepsilon_r$. Thus, this simple capacitor model predicts the monotonic increase in the permittivity as a function of h. Note, however, that Ballenegger and Hansen have questioned the layering picture [10].

As the dielectric response is anisotropic, the permittivity cannot be described by a single scalar, but must be considered as a general tensor property [30], in this case a so-called rank-2 tensor which will be introduced in Chapter 2. For some dielectric

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

materials, for example, multi-component crystals, the anisotropy can result in a polarisation which is not parallel to the electric field. To complicate the problem further, the preceding model also indicates that the permittivity is position dependent, that is, the permittivity tensor is a function of position in general.

It is not only the dielectric response which features anisotropy and position dependency [102]. Mechanical properties like viscosity also should, in principle, be considered to be anisotropic and position dependent. In our exploration we will see that this is complication is usually not needed for simple systems unless we study the dielectric properties.

1.1.4 Coupling Phenomena

In a molecular dynamics study, de Luca et al. [150] investigated water confined in a nanoscale slit-pore, where a rotational electric field was applied to the system; see the illustration in Fig. 1.5. The authors designed the slit-pore such that one wall was a graphene wall, a hydrophobic material, and the other wall was β-cristobalite which is hydrophilic. The water molecules' dipoles will align with the field; of course, due to thermal fluctuations this alignment is far from perfect. As the field rotates, the molecules will also rotate because the field exerts a torque on the dipoles and in this way the water molecules obtain a non-zero average angular momentum. Due to conservation of total angular momentum, this intrinsic molecular rotation results in a fluid flow, that is, a fluid translational motion. This coupling is not included in the classical hydrodynamic description, where the local rotation is given directly by the (local) curl of the streaming velocity field and is therefore not treated as an independent dynamical variable.

On a small historical note, the coupling was already described in the late 1890s by the Cosserat brothers [43, 44] and again treated in great detail in the 1950s to the 1980s [3, 52, 67, 190]. With the increasing interest in nanoscale hydrodynamics in the 2000s it is again the focus of many research groups [28, 72, 150].

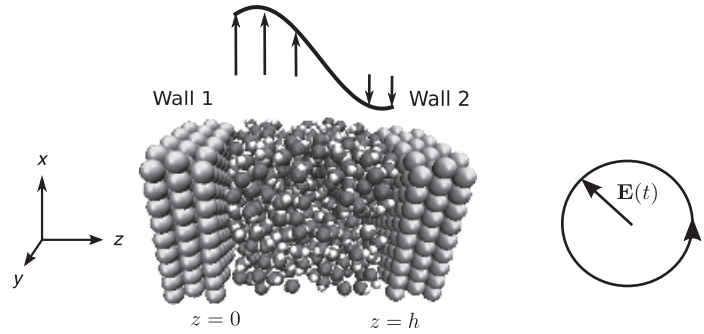

Figure 1.5 Illustration of the resulting velocity profile when a rotating electric field is applied to a slit-pore with water and non-symmetric wall hydrophobicity. Here wall 1 is hydrophobic, giving large slippage, and wall 2 is hydrophilic, giving small slippage. Figure from Ref. [102] with permission.

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

This particular coupling phenomenon has an interesting potential application. Pumping fluids that are confined in nanoscale geometries is a challenging task. Naively, one can apply a pressure difference, Δp, over the channel inlet and outlet. In a first approximation, assume that the volumetric flow rate is proportional to the pressure difference applied, $Q = \Delta p / R_{\text{hyd}}$, where R_{hyd} is the hydraulic resistance. It can be shown from classical hydrodynamics that $R_{\text{hyd}} \propto 1/h^4$ [38], and since h is in the order of nanometres, the hydraulic resistance is extremely large, leaving this simple pumping device unusable under normal circumstances. Other pumping mechanisms are therefore needed, and exploiting the coupling between the molecular rotation and the translational motion is one possibility. Also, Felderhof [71] proposes nano-propulsion systems based on this coupling.

The coupling between the molecular rotation and the flow is just one example of many coupling phenomena [52] relevant in nanoscale hydrodynamics. In nanoscale systems very large thermal gradients can be achieved, and these gradients can result in different mass fluxes for fluid mixtures, an effect referred to as the Soret effect. Also, Bresme et al. [35] showed that for polar fluids, a large thermal gradient induces polarisation. We return to these coupling phenomena in Chapter 6.

1.1.5 Non-local Viscous Response

The final example is from Todd et al. [199]. Here the authors investigated the fluid shear stress when applying a sinusoidal shear force to the fluid. The shear force acts in the x-direction and is a function of the z-direction; here we write it in terms of a force density $\rho g = \rho g_0 \cos(kz)$, where g_0 is the acceleration amplitude, and k defines wavelength of the imposed force; see Fig. 1.6. We assume that the system is homogeneous and infinite in extent, so we need not consider effects from confining walls. For this system there is only one non-zero shear stress component, σ_{zx}, where the first index indicates that the normal vector to the sheared virtual fluid surface is parallel to the z-direction, and the second index indicates the force direction; see Fig. 1.6. We here follow the original work and discuss the system response to the shear force through the stress; however, with a few exceptions, we use the shear pressure rather than the shear stress, as this is a more natural choice when deriving the momentum balance equation. Also, we will from here on omit writing the stress indices and simply use $\sigma = \sigma_{zx}$.

In the steady state, the momentum balance equation reads

$$\frac{\partial \sigma}{\partial z} = -\rho g = -\rho g_0 \cos(kz)\,. \tag{1.13}$$

Integration gives

$$\sigma(z) = -\frac{\rho g_0}{k} \sin(kz)\,, \tag{1.14}$$

using that the stress is zero at $z = 0$.

The stress can also be predicted using Newton's law of viscosity. For the geometry here we have

$$\sigma(z) = -2\eta_0 \dot{\gamma}(z)\,, \tag{1.15}$$

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

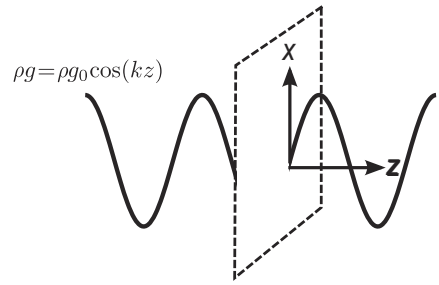

Figure 1.6 Illustration of the applied sinusoidal shear force and the fluid shear surface. The fluid is not shown.

where $\dot{\gamma}$ is the strain rate. For the applied force used here, the strain rate is a sine function, $\dot{\gamma} = \widetilde{\dot{\gamma}}(k)\sin(kz)/2$, hence,

$$\sigma(z) = -\eta_0 \widetilde{\dot{\gamma}}(k)\sin(kz)\,. \tag{1.16}$$

Both η_0 and the strain rate amplitude $\widetilde{\dot{\gamma}}$ can be found from independent methods.

The shear stress obtained from the momentum balance equation, Eq. (1.14), and the shear stress resulting from Newton's viscosity law, Eq. (1.16), are both shown in Fig. 1.7. The z-coordinate is given in units of around one atomic diameter, 3–4 Å. For small k-values (long wavelengths) Newton's viscosity law predicts the shear stress satisfactorily; however, as the wavelength approaches the atomic length scale (large k-values), the stress features significant reduction compared to the prediction from Eq. (1.16). It is worth noting that the force is sufficiently low such that the system is in the linear response regime [94].

The mechanism for this stress reduction is believed to reside in the non-local nature of the fluid response to the shearing force. Newton's viscosity law is local in the sense that the stress at some point, **r**, is proportional to the strain rate at that given point. More generally, the stress at **r** is dependent on the entire system strain rate distribution.

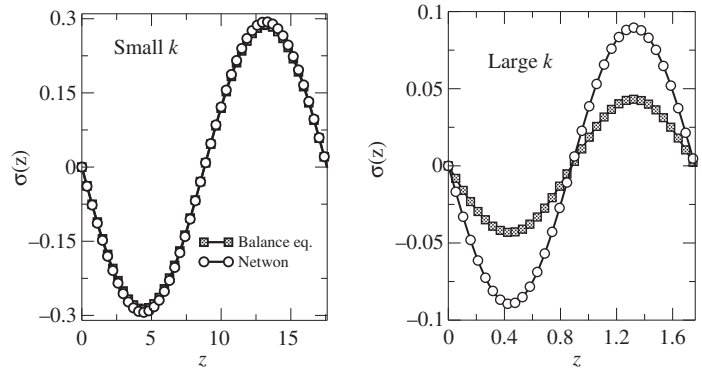

Figure 1.7 Shear stress profiles for an atomic fluid subjected to a sinusoidal shear force. In the small k-value figure the wavelength is one order of magnitude larger than in the large k-value figure. The z-coordinate is given in units of approximately one atomic diameter and the stress is in computer simulation units. Data redrawn from Ref. [102].

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

Phenomenologically, this can be modelled by letting the viscosity be a function of the distance between $\mathbf{r}$ and all other points in the system. That is, the viscosity is a non-local response function. This is equivalent to how we model the temporal viscoelastic response through a memory kernel; in the most general formulation the fluid response is modelled through time- and space-dependent, or equivalently frequency- and wavevector-dependent, response functions. This generalisation was initiated by Boltzmann [174] and is today referred to as generalised hydrodynamics; we return to this formalism in Chapter 4.

1.2 Nanoscale Hydrodynamic Modelling

How can we then model, that is better understand, these nanoscale hydrodynamic phenomena? The preceding discussions have all been based on hydrodynamic theory, that is, the continuum picture. However, we should be a bit concerned for at least two reasons. (i) One would think that the intrinsic discrete nature of the fluid at very small scales will destroy the continuum picture. This is also pointed out by Lautrup, who categorises continuum modelling as physics on the macroscopic scale [141]. (ii) At these small length scales the system may not behave according to classical mechanics, but be quantum mechanical in nature.

Let us first address the latter. Quantum mechanical effects become relevant when the system characteristic length scale is on the order of the de Broglie wavelength, λ_{Br}. The characteristic length scale is not always well defined; here we use the average distance between the molecules' centre of mass [90], which for liquids composed of small and approximately spherical molecules is around 3–4 Å. The de Broglie wavelength is [87]

$$\lambda_{\mathrm{Br}} \approx \frac{1}{\sqrt{mT}}\, 10^{-22}\mathrm{m}\sqrt{\mathrm{kg\ K}}, \tag{1.17}$$

where m is the molecular mass and T the temperature. Thus, for water $m = 2.99 \times 10^{-26}$ kg and at ambient temperatures $T = 300$ K, the de Broglie wavelength is 0.33 Å. Comparing this with the intermolecular distance, we must expect the quantum mechanical effects to be small. In fact, to dismiss the classical picture we will need to be in the very low temperature regime; say, for atomic hydrogen at the melting point $T = 13.95$ K we have $\lambda_{\mathrm{Br}} \approx 6$ Å, and here classical mechanics breaks down.

When dealing with molecules in the quantum mechanical realm, the molecular rotational energies are discretised, and we must also consider this effect. The different energy levels lead to a definition of a characteristic rotational temperature,

$$T_{\mathrm{rot}} \approx \frac{4 \times 10^{-46}\mathrm{kg\ m^2\ K}}{\Theta_{\mathrm{mol}}}, \tag{1.18}$$

where Θ_{mol} is the molecular moment of inertia. For small molecules, Θ_{mol} is on the order of $10^{-47} - 10^{-45}$ kg m^2 and the characteristic temperature typically fulfils $T_{\mathrm{rot}} <$ 50 K. Therefore, for the rotation we require that $T \gg T_{\mathrm{rot}}$ if quantum mechanical effects

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

are ignored, and this is the case for the systems we explore here. We will not treat molecular vibrational degrees of freedom in this text, and therefore assume that this set of dynamics has no effect on the phenomena we investigate.

What about molecular discreteness? In the continuum model we picture the fluid as being composed of small fluid elements. Each such element contains a sufficiently large amount of molecules such that one can define the same quantities (or properties) for the local fluid element as those of the fluid itself, no matter how small the fluid element is [207]. The quantities are here denoted hydrodynamic variables, or hydrodynamic quantities, and can be mass density, streaming velocity, energy, and so on.

Returning to the example of the fluid element $\mathcal{V}$ at the beginning of this chapter, we follow Lautrup [141] and consider the mass density, $\rho_{\mathcal{V}} = \rho_{\mathcal{V}}(t)$. The mass of $\mathcal{V}$ is simply the total mass of the constituent molecules inside $\mathcal{V}$, and $\rho_{\mathcal{V}}$ can then intuitively be written as

$$\rho_{\mathcal{V}}(t) = \frac{1}{\Delta\mathcal{V}} \sum_{i \,\text{in}\, \mathcal{V}} m_i \,, \tag{1.19}$$

where index i runs over all molecules contained in $\mathcal{V}$, m_i is the mass of molecule i, and $\Delta\mathcal{V}$ is the fixed fluid element volume. In the absence of a flow, $\rho_{\mathcal{V}}$ will fluctuate in time around the average density, ρ_{av}, as molecules enter and leave the fluid element due to thermal motion. The classical continuum model predicts a time-independent density $\rho_{\mathcal{V}}(t) = \rho_{\text{av}}$, or equivalently a time-independent number of particles N_{av} in $\mathcal{V}$. The fluctuations can then be thought of as a measure of the deviation between the continuum model and the actual situation.

Let us then quantify the fluctuations by the standard deviation, σ, around the average. If the molecule entering and leaving the fluid element is a true random event, then the standard deviation is proportional to $\sqrt{N_{\text{av}}}$. It is not the absolute standard deviation itself which is interesting here; one naturally expects the deviation to be smaller in the dilute case, as fewer molecules enter and leave the fluid element. Instead one can study the relative standard deviation σ/ρ_{av} which is proportional to σ/N_{av} or $1/\sqrt{N_{\text{av}}}$ when the volume $\Delta\mathcal{V}$ is fixed. Therefore, if we accept a relative error in the order of 1 per cent, we have that the average number of molecules in the fluid element must be greater than 10^4. There is no general rule for the acceptance threshold value; and this is perhaps not the interesting point here. The interesting point is that the relative standard error decreases with increasing density as $\sigma/\rho_{\text{av}} \propto 1/\sqrt{\rho_{\text{av}}}$. Thus, the continuum model performs poorly at small length scale for gasses when compared to more dense systems at the same length scale; also see Ref. [127].

To elaborate further, we study another very important hydrodynamic variable, namely the streaming velocity. This time we do not make the simple statistical arguments we have just made, but rely on a molecular dynamics simulation that includes correlation effects and so on; the molecular dynamics technique is introduced in what follows. In the first-order approximation we can ignore the density fluctuations [102], and the streaming velocity, $\mathbf{u}_{\mathcal{V}} = \mathbf{u}_{\mathcal{V}}(t)$, of $\mathcal{V}$ is intuitively given by

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

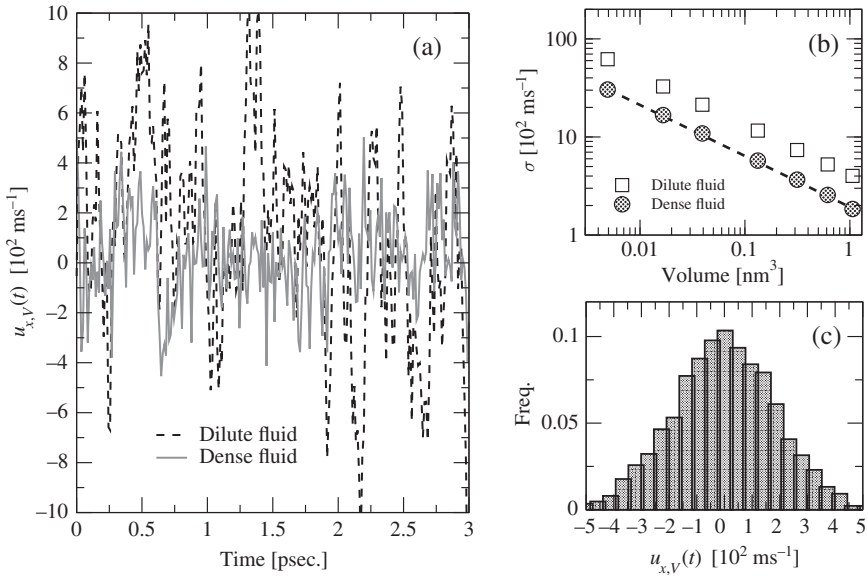

Figure 1.8 (a) Streaming velocity of a cubic fluid element with a volume of approximately 1 nm^3. (b) The corresponding standard deviation. (c) The normalised histogram for the streaming velocity x-component. The densities of the methane fluids are $\rho = 270$ kgm^{-3} and $\rho = 540$ kgm^{-3}, and the temperature is $T = 222$ K.

$$\mathbf{u}_{\mathcal{V}}(t) = \frac{1}{\rho_{\text{av}}\Delta\mathcal{V}} \sum_{i \text{ in } \mathcal{V}} m_i \mathbf{v}_i(t) \,, \tag{1.20}$$

where $\mathbf{v}_i$ is the velocity of molecule i. Note that the streaming velocity is defined from the molecular momenta and gives the correct mass weighted average. Again, there is no advection in the system, and we simply study the effect of thermal fluctuations. Figure 1.8(a) shows molecular dynamics simulation results for the x-component of the streaming velocity of a cubic fluid element with $\Delta\mathcal{V} \approx 1$ nm^3 for two cases: (i) a relative dilute methane fluid and (ii) a relative dense methane fluid. One can see that the fluctuations in the dilute case are larger than in the dense case. The underlying reason for this important result is that in the dense case the molecules collide, or more precisely interact, very frequently resulting in a high degree of momentum exchange compared to the dilute case. The statistics are summarised in Fig. 1.8(b), where the standard deviations around zero for the two state points are shown for different fluid element volumes. The punctured line shows a power-law function with an expected exponent of $-1/2$. If we accept a standard deviation threshold of around 200 ms^{-1}, that is, we demand 68 per cent of the data points to be within $u_{x,\mathcal{V}}(t) = (0 \pm 2) \times 10^2$ ms^{-1}, then the continuum model fails for characteristic length scales below 1 nm in the dense situation. This length scale increases as the density decreases in agreement with the preceding discussion of momentum exchange. In Fig. 1.8(c) the normalised frequency for $u_{x,\mathcal{V}}$ is plotted in the dense case; this plot indicates that the fluctuations are Gaussian distributed.

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

Strictly, the definitions of mass density and streaming velocity, Eqs. (1.19) and (1.20), make no sense in the coarse-grained continuum picture since there are no such things as molecules. They must therefore be thought of as microscopic (or molecular) definitions of the hydrodynamic variables, and not as continuum definitions.

If we insist on a *classical* continuum picture, the thermal fluctuations are indeed problematic. However, we could choose to include the fluctuations in our modelling and, inspired by the Langevin equation, one strategy that comes to mind is to simply add a stochastic noise (or force) term to the dynamical equations for the hydrodynamic variables. In this way the many degrees of freedom behind the thermal fluctuations are coarsened into a single random term, keeping the problem low-dimensional. Adding a fluctuating stochastic term onto the dynamics means adding energy to the system, and this must be balanced correctly by an energy drainage, or dissipation, originating from the system transport processes. Thus, the stochastic force amplitude cannot be chosen ad hoc, but depends on the system transport coefficients and the temperature; this is the so-called fluctuation-dissipation theorem. We will adopt the stochastic force method; however, we only explore situations where we do not need to invoke the fluctuation-dissipation theorem, and simply require that the force has the following properties:

1. the average over an ensemble is zero, and
2. it is uncorrelated with respect to the hydrodynamic variables.

The term 'ensemble' is here used in the general sense as a 'sufficiently large and statistically independent set' and not as a specific statistical mechanical ensemble. Adding such a stochastic force term to the dynamical equations for the hydrodynamics variables leads to a set of stochastic differential equations which pose new challenges. However, by performing an ensemble average over a set of independent initial conditions, the equations become deterministic due to the stochastic force properties we have just listed. For a more careful treatment and discussion of how to treat these hydrodynamic fluctuations, the reader is referred to the original work by Landau and Lifshift [139, 140] as well as the book by Zárate and Sengers [56].

By averaging we then suppress the thermal fluctuations, but this still does not answer our original question, namely whether the continuum picture can be applied to nanoscale fluid systems, or more precisely, whether the underlying physico-chemical processes can be modelled using continuum theory; this may fail even in the absence of fluctuations. We shall address this fundamental question in great detail throughout the book; but before doing so, we saw in the preceding examples one important point which is worth recalling:

In the continuum picture, the dynamics of the different hydrodynamic variables, say, the mass and momentum densities, are given through the balance equations. The balance equation is a partial differential equation; however, as such it does not form a mathematical closed problem that can be solved. Therefore, one applies constitutive equations (or relations) that typically relate the diffusive processes with the system gradients. The constitutive relations are models, and not fundamental laws of physics. Newton's law of viscosity, which we discussed for the shear stress, is an example of such a model. Then, the continuum description is based on the balance equation

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

and a set of constitutive relations. If the theory fails, it may not be because the continuum picture fails, but because the constitutive relations do not model the underlying physico-chemical processes appropriately. The non-local shear response is an example where the classical constitutive relation is not appropriate; however, by proposing a generalised constitutive relation the continuum picture is indeed applicable.

Thus, what may appear to be a breakdown of the continuum picture can simply be a result of poor modelling and lack of generality.

Often the next step in the theoretical treatment of nanoscale hydrodynamics is based on Mori–Zwanzig projection-operator formalism [161, 219] that leads to a generalised Langevin-type equation for the hydrodynamic variable. This equation depends on a function describing the transport properties, and this can be found from simulations or theoretically from mode coupling theory [81]. A related yet slightly different theory is the generalised collective mode theory [54], where the set of hydrodynamic variables is increased, leading to a larger dynamical space and consequently a better quantitative agreement between data and theory [40]. These advanced theories are successful in predicting important phenomenology; however, these are not the topic here, and while we will treat a lot of the same phenomena, we will do so from a purely hydrodynamic viewpoint.

1.3 Molecular Dynamics

The microscopic definitions of the hydrodynamic variables given in Eqs. (1.19) and (1.20) are based on the molecular positions and momenta. If a system is composed of N molecules, we can envision that a particular system state is defined by $3N$ position coordinates and $3N$ momentum coordinates; this is the molecular phase space point. As the molecular positions and momenta evolve in time, the phase space point changes, resulting in the system phase space trajectory. Molecular dynamics (MD) is a powerful simulation method to trace out this phase space, or at least parts of it and hopefully the important parts. From the dynamics of the phase space we can gain general knowledge of the fluid properties and in particular the hydrodynamics.

It is worth mentioning that many other microscopic and mesoscopic simulation methods have been applied to study small-length-scale hydrodynamics; these include Monte Carlo methods [5], the smooth particle applied mechanics (SPAM) method [111], lattice Boltzmann/lattice gas automata [42, 113], and the direct simulation Monte Carlo (DSMC) method [20], most of which are mainly applicable for gasses. Molecular dynamics is highly versatile and plays an ever increasing role in studying nanoscale phenomena [26], and we will apply only this simulation method. The following is not meant to be a thorough discussion of the molecular dynamics technique, but rather an introduction to the necessary terminology that will be used in the remaining chapters, and, importantly, also a justification for the use of molecular dynamics to explore nanoscale hydrodynamics. The interested reader is referred to the classical books on the subject, for example, Refs. [5, 74, 184].

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

Molecular dynamics is founded in classical mechanics, wherein Newton's second law is integrated numerically for each particle in the system. Here a particle can be a single atom, a molecule, or a group of atoms that move in a coherent fashion. If $\mathbf{F}_i$ is the force acting on particle i, having position $\mathbf{r}_i$, velocity $\mathbf{v}_i$, and mass m_i, the equations of motion are

$$\frac{\mathrm{d}\mathbf{r}_i}{\mathrm{d}t} = \mathbf{v}_i \text{ and } m_i \frac{\mathrm{d}\mathbf{v}_i}{\mathrm{d}t} = \mathbf{F}_i \,. \tag{1.21}$$

Then, for N particles we solve $6N$ coupled and, in general, nonlinear differential equations. In molecular dynamics we obviously assume that quantum effects can be safely ignored. As we have noticed, at ambient conditions this will indeed be the case for most liquids and fluids, as the de Broglie wavelength is less than one angstrom and significantly smaller than the relevant length scale. Not all phenomena can be described correctly by classical theories; for example, quantum mechanical effects occurring inside the wall may be relevant for the wall–fluid interactions and therefore also the hydrodynamics in highly confined geometries. These systems are not treated here.

Molecular dynamics relies on accurate models for the particle interactions; these inter-particle interactions are conservative and are therefore often given through a potential function U. In standard simulations we assume that the particles are spherical symmetric point masses and that the interactions are only pairwise such that $\mathbf{F}_i^c = -\boldsymbol{\nabla} U(r_{ij})$, where $\mathbf{F}_i^c$ is used to underline that the force is conservative and r_{ij} is the distance between particle i and particle j. The famous Lennard–Jones pair potential reads

$$U_{LJ}(r_{ij}) = 4\varepsilon \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - \left(\frac{\sigma}{r_{ij}} \right)^{6} \right] , \tag{1.22}$$

where ε and σ represent the interaction strength and characteristic diameter of the particles, respectively; most importantly here is that σ is typically on the order of a few angstroms. We are already now running out of symbols and do not want to confuse the length scale symbol with the standard deviation, or the energy scale with the dielectric constant. The first term is a repulsive term which accounts for the force due to electron repulsion at small inter-particle distances. The second term models the induced dipole moment; it is longer ranged and attractive. This is also known as the London dispersion force.

We will often use molecular dynamics simulations of methane fluid as controlled numerical experiments to test the hydrodynamic theories. The molecule is approximated to be a point mass spherical molecule – by far most of the mass is located in the carbon atom nucleus. The intermolecular interactions are modelled via the Lennard–Jones potential, where $m = 16$ g/mol, $\varepsilon/k_B = 148$ K, and $\sigma = 3.7$ Å [155]. The results from the molecular dynamics simulations can be presented in standard SI-units, but sometimes it is more convenient and insightful to list the results in units of σ, ε, and molecular mass m_i. Importantly, the unit of time in this unit system is $\sigma\sqrt{m/\varepsilon}$ and length is σ. We will use both SI units and molecular dynamics unit; the latter we denote MD units.

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

Table 1.1 Self-diffusivity, D_s, shear viscosity, η_0, relative permittivity, ε_r, heat of vaporisation, ΔH_{vap}, isothermal expansivity, χ_T, Debye relaxation time, τ_D, and thermal expansivity, α_p, for the SPC/Fw water model under ambient conditions. Both molecular dynamics results and the corresponding experimental values are listed. From Ref. [213].

	Units	SPC/Fw (MD)	Exp.
D_s	10^{-9} m^2 s^{-1}	2.35 ± 0.05	2.3
η_0	10^{-3} Pa s	0.75	0.85
ε_r		80 ± 2	78.5
ΔH_{vap}	kcal mol^{-1} K^{-1}	10.7 ± 0.1	10.52
χ_T	10^{-5} atm^{-1}	4.50	4.58
τ_D	psec.	9.5	8.3
α_p	10^{-4} K^{-1}	4.98	2.0

A lot of general information can be obtained by studying point mass particles like methane. However, some phenomena require that we use a more detailed model for the molecules, and today the molecular dynamics community simulates complex molecular systems with advanced interaction models; see Sadus [186] and Leach [142] for an overview. In general, we can write up a force field model that also includes Coulomb interactions, covalent bonds, forces due to angles and dihedral angles, and so forth. In terms of the potential function this is written as

$$U = U_{\text{LJ}} + U_{\text{coulomb}} + U_{\text{bonds}} + U_{\text{angles}} + U_{\text{dihedral}} + \ldots. \tag{1.23}$$

Widely used models are, for example, the CHARMM [36] and OPLS [122] force fields that give explicit expressions for the different terms and model parameter values depending on the specific systems under investigation. Again, we will not go into detail with the different interaction models, as this is far outside the scope of the text. It is, however, important to highlight that molecular dynamics can, with an accurate interaction model, predict the different mechanical, dynamical, and thermodynamic properties quite well under normal pressures and temperatures, where nanoscale fluid systems often operate. As an example of this, Table 1.1 lists different physical coefficients calculated from molecular dynamics simulations at equilibrium for the flexible simple point charge water model (SPC/Fw) [181, 213]. For comparison purposes, the corresponding experimental values are also given. Except for the thermal expansivity, the model results agree well with the experimental measured values. This indicates that molecular dynamics indeed can capture many of the underlying physical processes correctly, including the processes relevant for hydrodynamics. Importantly, water is not easily modelled, as the different properties of the liquid are a result of the complex long-ranged hydrogen-bond network.

Since we are solving a very large set of coupled differential equations, the number of particles and the time we can reach are both limited. In molecular dynamics large

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

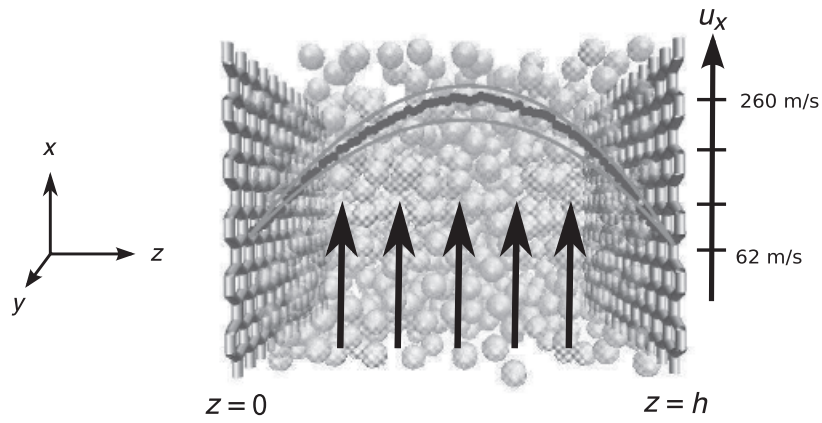

Figure 1.9 Molecular dynamics simulation of a planar Poiseuille flow. Symbols represent simulation data for the streaming velocity, and lines are the Navier–Stokes predictions; the two lines indicate the extremes in the uncertainty coming from the viscosity. Arrows illustrate the external force. Reprinted with permission from Ref. [102].

implies small: in 2013, the SuperMUC supercomputer simulated an impressively large system of 4.125×10^{12} particles, but here each time step (corresponding to approximately a femtosecond) took the computer 40 seconds, seriously limiting the time scale that can be studied. In the other extreme, one can reach 2×10^4 integration times steps per second for a small system size of 10^3 particles using Graphical Processor Units (GPUs) [9]. In the context of nanoscale hydrodynamics the number of particles is usually not too critical. For example, to simulate a methane fluid flow in a slit-pore geometry of height 10 nm we will need around 27×10^3 methane molecules if the simulation box is a perfect cube. For water in the same geometry we need approximately 10^5 hydrogen and oxygen atoms. The problem often lies in reaching realistic times, especially for charged systems like water, where the long-ranged Coulomb interactions are very computationally demanding. Even with a small number of particles, the time reached, τ_{obs}, with current computers usually does not exceed 10–100 nanoseconds. Thus, the phenomena we study with molecular dynamics must have small characteristic time scales. Often we must apply large external forces in order to excite the relevant physical mechanisms needed to reach a sufficiently small Deborah number, $\text{De} = \tau_s/\tau_{\text{obs}}$.

Figure 1.9 shows a snapshot from a molecular dynamics simulation of a methane fluid flowing in a slit-pore geometry similar to the one in Fig. 1.1. The molecular interaction is modelled through the Lennard–Jones potential, Eq. (1.22). The confining walls are also composed of Lennard–Jones particles and are positioned in a graphene-type lattice. The flow is generated by applying a constant external force that acts on the centre of mass of the methane molecules, and the resulting viscous heating is removed by applying a thermostat. In this manner a planar Poiseuille flow is simulated.

Even if equilibrium molecular dynamics results for the transport properties agree well with the experimental data, it is by no means trivial that the hydrodynamic model, in this case the Navier–Stokes equation, and the non-equilibrium simulations agree.

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

To make a direct comparison for the system in Fig. 1.9, we note that $h = 3.3$ nm, the density is $\rho = 270$ kgm^{-3}, the shear viscosity is $\eta_0 = 9.3 \pm 0.6$ μPa·s, and that the external acceleration applied is a staggering $g = 5.0 \times 10^{12}$ ms^{-1}. Using $U \approx 100$ ms^{-1} as the characteristic velocity, the Reynolds number, Eq. (1.2), is around 10. Hence, the flow is laminar. Furthermore, as the speed of sound for methane at this state point is approximately 10^3 ms^{-1}, the Mach number $M = U/c_s < 0.3$ and we need not consider fluid compressibility effects. In the slit-pore geometry the Navier–Stokes equation is reduced to a tractable boundary value problem

$$\eta_0 \frac{\mathrm{d}^2 u_x}{\mathrm{d}z^2} + \rho g = 0, \tag{1.24}$$

with

$$u_x(0) = u_x(h) = u_w = 62 \text{ ms}^{-1}, \tag{1.25}$$

where u_w is the slip velocity at the wall. The solution for this problem gives the well-known Poiseuille flow quadratic profile

$$u_x(z) = \frac{\rho g}{2\eta_0} z(h - z) + u_w \, . \tag{1.26}$$

The prediction from the Navier–Stokes equation is also shown in the figure (lines), and is in very good agreement with the time-averaged simulation data (filled circles). This is a crucial result; the hydrodynamic prediction, Eq. (1.26), where molecular details are strictly not considered agrees with the (time-averaged) molecular simulation results coming from solving the Newtonian equation of motion for each molecule, Eq. (1.21). The scenario that both of these very different descriptions of the flow are incorrect yet still produce the same result is highly unlikely. While definitely not being a proof, this example illustrates that hydrodynamics can for simple systems be applied on the nanoscale, *and* molecular dynamics can be applied to perform idealised numerical experiments of nanoscale fluid systems. Now, almost all real nanoscale fluid systems are not as simple and idealised as this, and laboratory experiments must, of course, always be the final test of our theoretical predictions.

As we have mentioned, the time scales we can reach with molecular dynamics are small compared to typical hydrodynamic time scales. In our Poiseuille flow simulation the external acceleration applied to drive the flow is of literally astronomical magnitude. This is necessary in order to obtain a well-developed velocity profile within the nanoscale time frame available. This large acceleration produces unrealistically large streaming velocities and strain rates, and yet, the simulation data agree with the Navier–Stokes predictions. The reason for this lies in the fact that the local Newtonian viscosity law for shear stress, Eq. (1.15), applies, that is, the system response is still linear and local. Lennard–Jones-type systems show a strain-rate-independent viscosity (Newtonian behaviour) for strain rates less than 10^{10}–10^{11} s^{-1} in the liquid phase [200] which is the same order of magnitude as the flow in Fig. 1.9. Non-Newtonian effects must always be considered, and for more complex fluids the regime where Newton's viscosity law is valid may not be accessible by molecular dynamics even when using highly optimised algorithms and hardware [146].

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

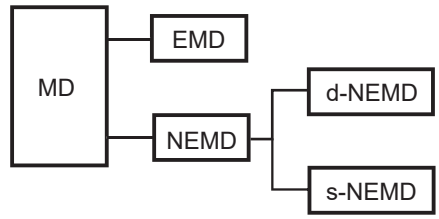

Figure 1.10 Schematic diagram of the three standard molecular dynamics (MD) techniques. Equilibrium (EMD), direct non-equilibrium (d-NEMD), and synthetic non-equilibrium (s-NEMD) simulations.

1.3.1 Molecular Dynamics Simulation Techniques

Simulations in molecular dynamics are divided into two main categories, equilibrium simulations (EMD) and non-equilibrium simulations (NEMD). With EMD we simulate the system in some well-known statistical mechanical ensemble, for example, the microcanonical ensemble where the number of molecules, volume, and energy are constants along the phase space trajectory. In this way we can use classical statistical mechanical results for the particular ensemble to derive the given properties we are studying. The properties listed in Table 1.1 are an example of a simulation in the canonical ensemble, where the number of molecules, volume, and temperature are constants. In Chapters 3 and 4 we rely heavily on equilibrium simulations to test how hydrodynamics predict relaxation phenomena in equilibrium.

In Chapters 5 and 6 we explore non-equilibrium systems. In non-equilibrium we can perform either direct (d-NEMD) or synthetic (s-NEMD) simulations. For d-NEMD we try to mimic the real physical experiment, at least to some approximation. Our Poiseuille flow is an example of this. The fluid flow is generated by application of some external driving force, $\mathbf{F}_i^{\text{ext}} = m_i\mathbf{g}$, by moving a wall or similar, and the resulting viscous heating is removed by thermostating the confining wall atoms; a discussion of the different thermostating methods can be found in Ref. [18]. The d-NEMD method is relatively straightforward to implement; for example, the driving force can be added directly to Eq. (1.21), so that for the fluid particles we have

$$\frac{\mathrm{d}\mathbf{r}_i}{\mathrm{d}t} = \mathbf{v}_i \text{ and } m_i\frac{\mathrm{d}\mathbf{v}_i}{\mathrm{d}t} = \mathbf{F}_i = \mathbf{F}_i^{\text{ext}} + \mathbf{F}_i^{c}\,. \tag{1.27}$$

However, d-NEMD is not always suitable if one needs to study and isolate a specific fluid phenomenon, as the confining walls often complicate and clutter the problem. Another problem with d-NEMD is that the statistical mechanics for such systems is not well developed, and we need to approach the analysis of our simulation results with care.

To overcome the problem associated with d-NEMD, one can perform s-NEMD simulations. Here the equations of motion are changed in order to probe a specific dynamical feature. The particles do not follow the simple Newtonian equations

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

of motion, but their dynamics are changed synthetically such that the system possesses, for example, a constant [65, 145] or spatially varying [12, 80] strain rate while keeping the local density and temperature constant (on average). s-NEMD is for this reason also referred to as homogeneous NEMD. In Section 1.1.5 this technique was used to investigate the fluid response to an imposed sinusoidal strain rate.

It is important to note that while the equations of motion are synthetic (also referred to as fictitious) and often not realisable in the laboratory, the system trajectory follows Gauss' principle of least constraints [66], and we can expect that the s-NEMD method probes the correct physics. The hydrodynamic equations derived from s-NEMD equations of motion are effected by the thermostat and are not the same as those derived directly from Eq. (1.21). Therefore, great care must be taken when interpreting and analysing results from such simulations [200]. Again, it is outside the scope of this text to pursue an in-depth introduction to the different MD techniques. For more details on how to implement the d-NEMD and s-NEMD methods, the reader is referred to the books by Evans and Morriss [66] and Todd and Daivis [200].

1.3.2 Mesoscale Molecular Dynamics

In order to extend the time scales we can reach with standard molecular dynamics methods, that is, increase τ_{obs} and reduce the Deborah number, alternative simulation methods have been devised. One such method is dissipative particle dynamics (DPD) [63, 112]. Rather than solving Newton's equations of motion for the individual atom or molecule, DPD solves the equation of motion for a collection of particles moving in a coherent fashion. This coherent motion is described through a single DPD particle, and in order to account for the coarse graining, the random force, $\mathbf{F}_i^R$, and dissipative force, $\mathbf{F}_i^D$ terms are augmented to the Newtonian equation of motion, that is,

$$\frac{\mathrm{d}\mathbf{r}_i}{\mathrm{d}t} = \mathbf{v}_i \text{ and } m_i \frac{\mathrm{d}\mathbf{v}_i}{\mathrm{d}t} = \mathbf{F}_i^c + \mathbf{F}_i^R + \mathbf{F}_i^D, \tag{1.28}$$

in the absence of any external forces. Importantly, these two forces are defined such that the total momentum is conserved, and this makes DPD fundamentally different from Brownian simulations and ensures hydrodynamic conservation of momentum. As with standard molecular dynamics, the force $\mathbf{F}_i^c$ represents the conservative interactions with other particles; but due to the coarse graining, these interactions are usually modelled as being 'soft', allowing the DPD particles to overlap; and a larger integration time step can be applied, which is important if we wish to simulate the system for a larger period of time. DPD is then based on a set of stochastic differential equations and is a mesoscopic description of the fluid. It is not always straightforward to extract the physical time and length scales in such simulations, that is to say, what size a DPD particle has. Moreover, while it has been shown that the balance equations for mass

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press

and momentum densities are obeyed, the total energy density is not conserved due to the random and dissipative forces [154]. However, it has been shown that DPD do capture many of the underlying hydrodynamics processes [103]; it is therefore a very potent alternative to classical molecular dynamics simulations, and we will use DPD to explore viscoelastic phenomena.

https://doi.org/10.1017/9781009158749.002 Published online by Cambridge University Press