University of Strathclyde

Department of Mechanical & Aerospace Engineering

Rarefied gas flows in microscale geometries: a hybrid simulation method

Stephanie Y. Docherty

A thesis presented in fulfilment of the requirements
for the degree of Doctor of Philosophy

2015

Declaration of author's rights

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

Stephanie Y. Docherty

June 2015

Abstract

Accurate predictions of the flow behaviour in microscale geometries are needed, for example, to design and optimise micro devices, and to ensure their safety/reliability. Rarefied gas flows in such geometries tend, however, to be far from local thermodynamic equilibrium, meaning that the flow behaviour cannot be described by conventional fluid mechanics. Alternative approaches for modelling 'non-equilibrium' gas flows have been proposed in recent years; because analytical solution methods are subject to significant limitations, the direct simulation Monte Carlo (DSMC) method is, at present, the most practical numerical simulation tool for dilute gases.

Unfortunately, the computational expense of tracking and computing collisions between thousands (or perhaps millions) of DSMC particles means that simulating the scales of realistic flow problems can require months (or even years) of computing time. This has resulted in the development of continuum-DSMC 'hybrid' methods, which aim to combine the efficiency of a conventional continuum-fluid description with the detail and accuracy of the DSMC method.

This thesis focuses on the development of a continuum-DSMC method that offers a more general approach than existing methods. Using a heterogeneous framework with a field-wise coupling strategy, this new method is not subject to the limitations of the well-known domain decomposition framework, or the restrictions of the heterogeneous point-wise coupling approach. The continuum-fluid description is applied across the entire flowfield, while the DSMC method is performed in dispersed micro elements that can be any size and at any location; these elements then provide the continuum description with updated constitutive and boundary information. Unlike most methods in the literature, the coupling strategy presented here is able to cope with heat transfer, and so non-isothermal flows can be simulated.

Testing and validation of this new continuum-DSMC method is performed by simulating a number of benchmark cases and comparing the results with full DSMC solutions of the same cases. Two 1D flow problems are considered: a micro Fourier flow problem tests the energy coupling procedure of the method, and a high-speed micro Couette flow problem demonstrates the full coupling algorithm. In general, the method's accuracy is found to depend on the arrangement of the micro elements — with sufficient micro resolution, good agreement with the equivalent full DSMC simulations can be obtained. Although the hybrid method offers no computational speed-up over the full DSMC simulations for several of these 1D test cases and only modest speed-ups for the others, both of these 1D flow problems are simulated only to validate the coupling strategy of the method.

Considerable speed-ups are offered by the method when simulating a larger and more realistic flow problem: a microchannel with a high-aspect-ratio cross-section acts as a representative geometry for modelling a gas flow through a narrow microscale crack. While the limitations of existing hybrid methods preclude their use for this type of high-aspect-ratio geometry, the new hybrid method is able to model this problem under isothermal and non-isothermal conditions. The implementation of the method is simplified to 2D by assuming that the flow variation in the streamwise direction is negligible, i.e. the method is applied to the microchannel cross-section only. Accurate predictions of the mass flow rate and the streamwise velocity field are obtained for a number of test cases; accurate predictions of the temperature field are also obtained when there is a temperature difference between the bounding walls.

Acknowledgements

First I want to express my sincere gratitude to my supervisor, Prof. Jason Reese. His guidance, advice, and continuous support has been invaluable, and his success sets an example that I can only hope to emulate in the future.

I am also sincerely grateful to a number of others who helped make this thesis possible. I am truly indebted to Dr Duncan Lockerby and Dr Matthew Borg for sharing their insights into multiscale simulation methods, and for many useful and enlightening discussions. I am grateful to Dr Craig White for answering many queries on the DSMC method and rarefied gas dynamics, and for sharing his expertise in programming.

I would like to thank Dr William Nicholls for helping me learn to use OpenFOAM, and Alexander Patronis for his help with MATLAB. I would also like to thank all of the researchers that are involved in the 'Micro and Nano Flows for Engineering' research partnership and everyone in the James Weir Fluids Laboratory for many informal discussions.

I am grateful for the financial support of the EPSRC, which has allowed me carry out this work.

Finally, I would like to thank my family (especially my mum and dad) and my closest friends. This work would not have been possible without their unconditional love, support, and encouragement. I dedicate this thesis to them.

Contents

List of Figures

List of Tables

Nomenclature

A	area [m^2]
$\mathbf{a}$	acceleration vector [m/s^2]
b	bin index [-]
$\mathbf{c}$	instantaneous molecular/particle velocity vector [m/s]
$\mathbf{c}'$	particle thermal velocity vector [m/s]
c'	particle thermal speed [m/s]
$\mathbf{c}_m$	centre of mass velocity for two colliding particles [m/s]
$\mathbf{c}_r$	relative velocity between two colliding particles [m/s]
c_r	relative speed between two colliding particles [m/s]
d	molecular diameter [m]
e	specific energy [J/kg]
$\mathbf{F}$	external body force vector [N]
$\mathbf{F}_{ij}$	force acting between molecules i and j [N]
F_N	number of real molecules that each DSMC particle represents [-]
$\mathbf{f}$	external body force vector per unit volume [N/m^{-3}]
f	normalised velocity distribution function [-]
f_o	Maxwellian distribution function [-]
G	number of DSMC time-steps [-]
G_{av}	number of DSMC time steps in an averaging period [-]
g	time step index [-]
H	vertical separation between walls/channel height [m]
$H_{\text{lower,upper}}$	vertical extents of near-wall micro elements [m]

H_{bulk}	vertical extents of bulk micro elements [m]
h	micro element index [-]
I	number of iterations [-]
i, j	particle or macro node indexes [-]
J	non-linear collision integral [$\text{m}^{-3}\text{s}^{-1}$]
Kn	Knudsen number [-]
k	iteration index [-]
k_B	Boltzmann's constant [J/K]
l	characteristic dimension [m]
M_{cand}	number of candidate pairs [-]
M_x	number of macro nodes across x dimension [-]
M_y	number of macro nodes across y dimension [-]
Ma	Mach number [-]
m	molecular mass [kg]
$\dot{m}$	mass flow rate [kg/s]
N	number of DSMC particles [-]
N_{mol}	number of molecules [-]
N_{macro}	number of implicit steps of a steady macro solver [-]
n	number density [m^{-3}]
n_p	order of a polynomial curve [-]
n_s	order of a B-spline curve [-]
$\mathbf{p}$	pressure tensor [N/m^2]
p	scalar pressure [N/m^2]
$\mathbf{q}$	heat flux vector [W/m^2]
R	specific gas constant [J/kgK]
R_f	random fraction [-]
$\mathbf{r}$	molecule/particle position vector [m]
S	computational speed-up [-]
s	property index [-]
T	temperature [K]
$T_{\text{cold,hot}}$	temperature of cold and hot walls [K]

t	time [s]
t_{cl}	average clock time per DSMC time step [s]
t_{mc}	mean collision time [s]
t_{micro}	relaxation period of a micro solver [s]
$\mathbf{u}$ $(= \mathbf{c}_0)$	macroscopic (stream) velocity vector [m/s]
u, v, w	components of the macroscopic velocity vector $\mathbf{u}$ [m/s]
$u_{\text{lower,upper}}$	horizontal speed of lower and upper walls [m/s]
V	volume [m^3]
W	horizontal separation between walls/channel width [m]
$W_{\text{left,right}}$	horizontal extents of near-wall micro elements [m]
W_{bulk}	horizontal extents of bulk micro elements [m]

Greek

α	diffuse fraction or TMAC [-]
Γ	correction field index [-]
γ	ratio of specific heats [-]
ΔN	number of DSMC particles to be inserted/deleted [-]
ΔT	temperature difference [K]
Δt	'macro' continuum time step [s]
Δx	horizontal spacing between macro nodes [m]
Δy	vertical spacing between macro nodes [m]
δ	mean molecular separation [m]
δp	difference between gauge and actual pressures [N/m^2]
δt	'micro' DSMC time step [s]
δx	bin width [m]
δy	bin height [m]
$\bar{\epsilon}$	mean percentage error [%]
ε	energy [J]
ζ	convergence parameter [-]
ζ_{tol}	convergence tolerance value [-]

Θ_{inc}	angle of incidence [°]
Θ_{refl}	angle of reflection [°]
ϑ	post-collision elevation angle [°]
κ	thermal conductivity [W/mK]
Λ	solid angle [sr]
λ	molecular mean free path [m]
μ	dynamic viscosity [N/m^{2}s]
ξ	particle collision cross-section [m^2]
Π	number of micro elements [-]
ρ	mass density [kg/m^3]
ς	relaxation constant [J]
$\boldsymbol{\tau}$	viscous stress tensor [N/m^2]
Υ	overall energy correction term [W/m^3]
$\boldsymbol{\Phi}$	heat flux correction vector [W/m^2]
φ	post-collision azimuthal angle [°]
Ψ	overall momentum correction term [N/m^3]
$\boldsymbol{\Omega}$	stress correction tensor [N/m^2]
ω	viscosity exponent [-]

Superscripts

$*$	post-collision properties

Subscripts

0	stream properties
av	averaged properties
c	properties of a cell
corr	corrected properties from the hybrid method
Full	properties of a full-scale DSMC simulation

gl	global properties
gs	properties of gas at wall surface
Hyb	properties of a hybrid simulation
inc	incident
l	local properties
max	maximum
meas	measured properties (before control)
min	minimum
NSF	Navier-Stokes-Fourier property fields
n	normal component
new	new properties (after control)
ref	reference properties
refl	reflected
RR	relaxation region
sur	surface properties
SR	sampling region
t	tangential component
tar	target properties (for control)
tot	total
tr	translational properties
V	properties of a volume
VHS	variable hard sphere properties
wall	wall properties
x, y, z	directional components of vectors

Acronyms

BGK	Bhatnagar-Gross-Krook
CA	Continuous micro solution – Asynchronous coupling
CAI	Continuous micro solution – Asynchronous Intermittent coupling
CE	Chapman-Enskog
CFD	Computational Fluid Dynamics
CI	Continuous micro solution – Intermittent coupling
CLL	Cercignani-Lampis-Lord
CM^3	Coupled Multiscale Multiphysics Method
DD	domain decomposition
DSMC	direct simulation Monte Carlo
DVM	discrete velocity method
G13	Grad's 13 moment equations
GPU	Graphic Processing Unit
HI	Heterogeneous micro solution – Intermittent coupling
HMM	Heterogeneous Multiscale Method
HMM-FWC	HMM with field-wise coupling
IMM	Internal-flow Multiscale Method
LJ	Lennard-Jones
MD	molecular dynamics
MEMS	microelectromechanical systems
MOB	method of bins
MOZ	method of zones
NS/NSF	Navier-Stokes/Navier-Stokes-Fourier
NTC	no time counter
OpenFOAM	Open Field Operation And Manipulation
R13	regularised 13 moment equations
SHMM	seamless HMM
TMAC	tangential momentum accommodation coefficient
VHS	variable hard sphere

Chapter 1

Introduction

1.1 Motivation

Conventional fluid mechanics is able to predict the flow behaviour in most fluid flow problems by assuming a continuum-fluid description. There are, however, situations where the flow can be far from local thermodynamic equilibrium and the molecular nature of the fluid (which is neglected in conventional fluid mechanics) becomes important. In this thesis, our interest lies in the behaviour of dilute gases; dense fluids behave somewhat differently. As illustrated in Fig. 1.1, 'non-equilibrium' flow (also commonly referred to as 'rarefied' flow) may occur in a gas when the density is low, when surface effects become dominant in highly confined geometries, or when there are large gradients in the flow properties. The conventional linear constitutive relations and the traditional 'no velocity slip/temperature jump' fluid boundary conditions are not accurate for non-equilibrium flows. Note that, in conditions of extreme non-equilibrium, the flow may become 'non-continuum' i.e. the continuous fluid description breaks down completely.

Until the late 1950s, non-equilibrium effects were generally found in low density gas flows, for example around high altitude aircraft or spacecraft. However, the trend towards miniaturisation within science and technology (popularised by Nobel Prize winner Richard Feynman in 1959 [1]) has led to the development of revolutionary microelectromechanical systems (MEMS) and nano-scale systems that display non-equilibrium flow behaviour. While MEMS technologies (i.e. devices with length scales between

100 nm and 100 μm) are readily available on the market today, nanotechnologies (i.e. systems with length scales in the range of 1 to 100 nm) are still emerging.

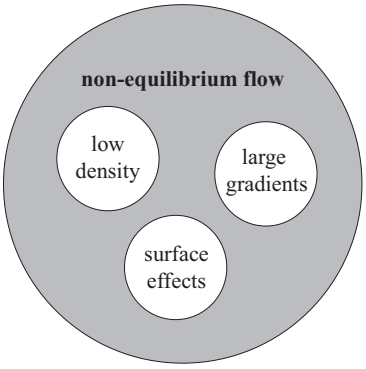

Figure 1.1: Causes of non-equilibrium flow in gases.

Due to their interdisciplinary applications, it is expected that micro- and nano-scale systems could have a major impact on many industries in coming decades [2]. The downscaling of systems can increase performance; examples include micro heat exchangers that can provide increased heat transfer for circuit cooling within the electronics industry; lightweight micro components in the aerospace industry that enhance flight performance and reduce material consumption; 'lab-on-a-chip' devices that enable faster (and potentially more accurate) medical diagnosis [3]; micro fuel cells and micro gas turbines that not only reduce component volume, but are also more powerful and longer lasting than conventional batteries; and carbon nanotubes, which can provide much greater mass flow rates than predicted by continuum fluid mechanics as well as excellent salt rejection capabilities for the desalination of salt water [4].

Despite advances in fabrication, the non-conventional flow behaviour through and around micro- and nano-scale devices is still not well understood. Such behaviour has been observed experimentally, for example, in an investigation performed by Arkilic *et al.* [5]: the mass flow rates of a rarefied gas inside long rectangular microchannels were found to deviate from that predicted by the 'no-slip' Navier-Stokes model, as shown in Fig. 1.2. Note that, while the first-order slip model is shown to match fairly well to the experimental values, the rarefaction is only slight with an outlet Knudsen number $Kn = 0.155$ (Kn provides an indication of the level of non-equilibrium in a gas, as will be discussed in Chapter 2). Later experiments by Colin [6] in similar microchannel

geometries confirmed that the accuracy of the first- and second-order slip models was reduced at higher Kn; see Fig. 1.3.

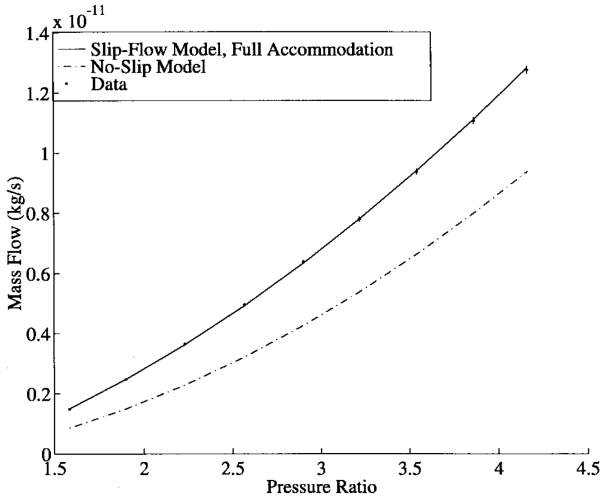

Figure 1.2: Taken from Arkilic *et al.* [5]: the mass flow rate of helium gas through a microchannel of height 1.33 μm with an outlet Knudsen number Kn=0.155. Dots indicate experimental values; dashed-dot curve is the no-slip solution; solid curve is the first-order slip solution.

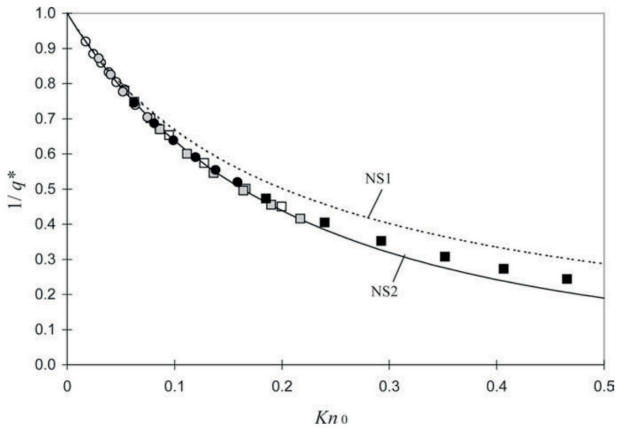

Figure 1.3: Taken from Colin [6]: inverse reduced flow-rate of gas through a microchannel for a pressure-ratio of 1.8. Curves NS1 and NS2 represent predictions from the first- and second-order slip models, respectively. Circles and squares indicate experimental values for nitrogen and helium gas, respectively; different colours (white, grey, and black) represent different channel geometries.

Poor performance of these slip models at high Kn was also shown by Lockerby and Reese [7] in the study of low-speed flow past a micro-sphere. By measuring the normalized drag, the new model that was presented (a new slip model with a near-wall scaling of the Navier-Stokes constitutive relations) was compared with the first- and second-order slip models, as well as Bhatnagar-Gross-Krook (BGK) results from

Lea and Loyalka [8], and a curve fit to the experimental data of Allen and Raabe [9]. As shown in Fig. 1.4, the new model and the BGK approach were significantly more accurate than even the second-order slip model at high Kn.

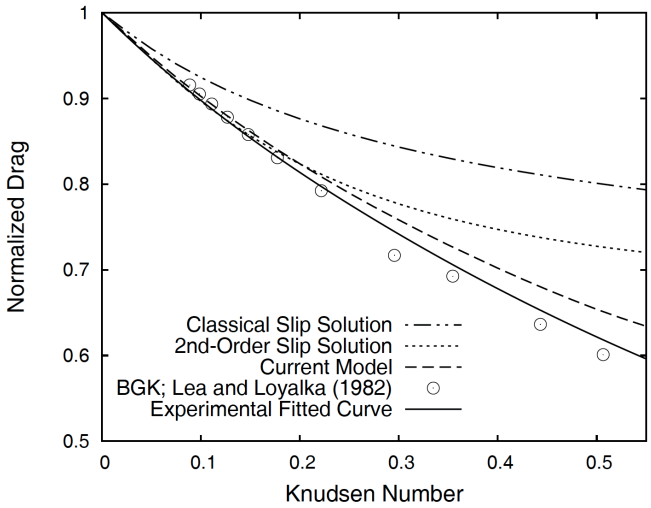

Figure 1.4: Taken from Lockerby and Reese [7]: normalized drag on a sphere. Comparison of classical slip solution; a second-order slip solution; new model; a BGK solution by Lea and Loyalka [8]; and a curve fit to the experimental data of Allen and Raabe [9].

Accurate flow prediction is needed for the design and optimisation of micro-devices, and so a greater understanding of the fundamental flow behaviour is essential. At the moment, there are significant limitations on the data that can be extracted from physical experiments due to the lack of suitable measurement techniques. Attention has therefore turned to modelling approaches that, unlike conventional fluid solvers, are able to capture the molecular behaviour of the fluid. Kinetic theory focuses on modelling the molecular behaviour of dilute gases by solving the Boltzmann equation. Generally, both analytical and numerical solution methods (for example, based on the BGK model) are subject to limitations; numerical simulation tools are, however, fairly robust and can be applied to a wide range of flow problems.

The most fundamental molecular simulation tool is molecular dynamics (MD) [10]. Taking a deterministic approach, this technique solves Newton's equations of motion for a system of interacting atoms/molecules. While being suitable for liquids and dense gases, MD is inefficient for dilute gases (where molecular interactions are less frequent). The direct simulation Monte Carlo (DSMC) method [11] can instead provide a coarse-

grained molecular description. Founded on kinetic theory, the DSMC method reduces computational expense by adopting two approximations: (1) each simulated particle can represent a large number of real gas molecules, and (2) the particle motion and interparticle collisions can be decoupled over small time intervals; particle movements are computed deterministically, while interparticle collisions are treated statistically.

Despite their accuracy, these molecular (or particle) simulation tools are generally much too expensive to resolve the spatial and temporal scales of real engineering problems. Multiscale 'hybrid' methods that combine the efficiency of continuum methods with the detail and accuracy of molecular descriptions have therefore been developed over the past few decades. The molecular treatment is applied over micro/nano scales to resolve the microscopic flow behaviour, while a continuum-fluid description is employed over macro scales to resolve macroscopic flow variations. The essential aim of these hybrid methods is to obtain computational savings by exploiting spatial and/or temporal 'scale separation' as much as possible; scale separation exists when there is only a loose coupling between the microscopic and macroscopic flow behaviour.

1.2 Objective: a continuum-DSMC hybrid method

The vast majority of existing continuum-DSMC hybrid methods are based on a domain decomposition (DD) framework [12–25]. The main reason for this is that the DD approach is able to exploit the fact that non-equilibrium flow often appears close to bounding walls: the molecular/particle solver is applied in a 'micro' sub-domain close to the wall, and a conventional continuum-fluid solver is applied in the remaining 'macro' sub-domain; an interface or overlap then enables coupling of the two solvers, as shown in Fig. 1.5(a). However, DD methods are limited to flow problems where the continuum-fluid model is accurate across the majority of the flowfield and microscopic resolution is required only in local regions (i.e. close to bounding walls or around singularities).

Although newer and less common than the DD framework, the Heterogeneous Multiscale Method (HMM) [26] has the ability to simulate problems that require the particle solver to complement the continuum solver everywhere, i.e. when the traditional linear fluid-constitutive relations fail within the bulk of a system. This could be the case if non-equilibrium flow appears in the bulk (caused, for example, by large temperature

gradients) or if the transport properties are unknown in an unusual gas mixture. In the HMM framework, a continuum-fluid description is applied across the entire flowfield, and the molecular/particle solver is applied in spatially-distributed micro elements. These micro elements provide the local continuum description with updated boundary conditions and/or fluid-constitutive information. As illustrated in Fig. 1.5(b), a 'point-wise coupling' approach is typically adopted — the micro elements supply this information directly to the nodes of the macro grid, both at the bounding walls and in the bulk. At the same time, the elements are constrained based on the local continuum solution obtained at the collocated macro nodes. Unfortunately, point-wise coupling means that the position and size of the micro elements is restricted by the placement and the density of the macro nodes. It also means that, while the HMM is effective when the spatial scales are highly separated (i.e. when the variation of the macroscopic properties relative to the physical extent of a micro element is small), it is inefficient for flow problems that possess mixed degrees of spatial scale separation.

To overcome the restrictions of point-wise coupling, the HMM with field-wise coupling (HMM-FWC) [27] was proposed fairly recently. Once again, this framework places a continuum-fluid description across the entire flowfield. However, each micro element now represents fields that correlate directly with an identically-sized continuum subregion. The coupling is then performed via fields rather than the nodal points: local continuum property fields are imposed across the micro elements, and local constitutive correction fields (i.e. corrections to the linear constitutive-fluid relations) are extracted; these local corrections are then interpolated to provide a global correction across the entire flowfield. Like in the HMM, the near-wall micro elements are also able to provide the continuum description with updated boundary information. Essentially, the HMM-FWC can be considered a more general heterogeneous approach than the HMM — as indicated in Fig. 1.5(c), the position and size of the micro elements is not restricted by the macro nodes, and flows with mixed degrees of spatial scale separation can be simulated.

In theory, the HMM-FWC could be applied to simulate the flow through long micro/nanochannels; such channels are fairly common in emerging micro- and nanotechnologies and so these high-aspect-ratio flows are of increasing interest. As a large

portion (if not all) of the flowfield is near-wall, the application of a DD method would be inefficient. The traditional HMM would also be unsuitable as the micro elements would likely be forced to overlap. However, a specific class of HMM-FWC has been developed for these types of flows: the framework and coupling strategy of the Internal-flow Multiscale Method (IMM) [28] is tailored to exploit the large length scale separation that exists in the streamwise direction. As in the other heterogeneous frameworks, the IMM applies a continuum-fluid description over the entire channel. Very short micro elements that occupy the entire cross-section (i.e. the entire channel height) are then distributed along the channel length; see Fig. 1.5(d). The coupling is simplified compared with other heterogeneous methods by circumventing the need to supply boundary conditions or constitutive information directly; pressure gradients are imposed across the micro elements and the resulting mass flux is used to correct the continuum description.

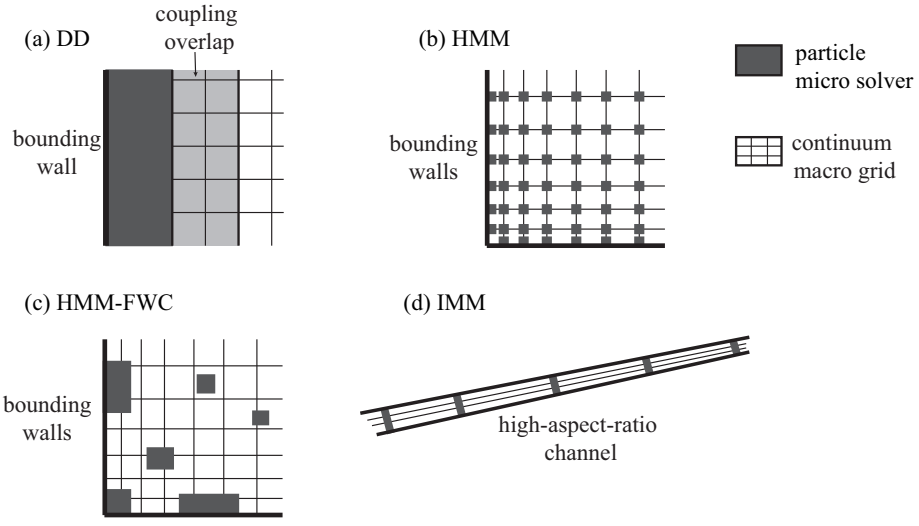

Figure 1.5: Schematic of the four hybrid frameworks: (a) domain decomposition (DD), (b) the Heterogeneous Multiscale Method (HMM), (c) the HMM with field-wise coupling (HMM-FWC), and (d) the Internal-flow Multiscale Method (IMM).

The HMM, the HMM-FWC, and the IMM were all originally implemented with a continuum-MD coupling for the simulation of liquids. With the exception of the IMM, the use of heterogeneous methods has continued to focus on continuum-MD couplings — there has been little attention given to continuum-DSMC couplings for dilute gases. Although the Coupled Multiscale Multiphysics Method (CM^3) of Kessler *et al.* [29] was

based on a heterogeneous approach (i.e. the DSMC method supplied the continuum-fluid description with constitutive information), both the continuum description and the DSMC solver were employed over the entire flowfield. The spatial and temporal scales remained fully coupled, and computational savings were obtained only by using a smaller number of averaging ensembles in the DSMC stage of the CM^3 than was used in the equivalent full DSMC simulations.

With the aim of a versatile hybrid method that can deal with mixed levels of spatial scale separation across a wide range of dilute gas flows, this thesis will adapt the HMM-FWC approach to a continuum-DSMC coupling.

Both the HMM-FWC and the IMM were first developed with assumptions of steady, incompressible, and isothermal flow. The IMM has now been implemented for a continuum-DSMC coupling [30, 31] and has seen extensions to compressible [30], non-isothermal [31], and unsteady flows [32]; the HMM-FWC has seen no such developments to date. Therefore, as well as modifying the constraint of the micro elements to suit the use of a DSMC solver, the hybrid method of this thesis will extend the HMM-FWC coupling strategy to include capabilities in compressible and non-isothermal flows. An extension to unsteady flows will not be considered, but could be a subject of future work.

With its flexibility, this new continuum-DSMC method will have the potential to simulate a wide range of flow problems: it will be applicable for the types of problems that are normally tackled by DD methods, i.e. where non-equilibrium appears at bounding walls; it will also be suitable when non-equilibrium appears in the bulk, or when the transport properties of the gas are unknown.

To demonstrate the method's ability to deal efficiently with mixed levels of spatial scale separation, it will be used to simulate a rarefied gas flowing in a microchannel that is representative of a microscale crack. Leakage from narrow cracks can occur, for example, in pressure vessels, pipes, and valves. The limitations of the other hybrid frameworks preclude their use for crack-type geometries where the cross-section of the flow is of high-aspect-ratio: DD methods would be inefficient as the entire flowfield is near-wall; the micro elements of the HMM would overlap; and, although the IMM could exploit high scale separation in the streamwise direction, it would require each

micro element to occupy the entire cross-sectional area and so would not be able to exploit scale separation over the large dimension of the flow cross-section. Our new HMM-FWC method will have the ability to exploit high scale separation over this dimension, while also dealing with the low scale separation that exists over the small dimension of the cross-section. Note also that, while the original HMM-FWC method was demonstrated for a 1D flow problem only, our new method will be implemented in 2D for this crack flow problem.

1.3 Thesis outline

Chapter 2 introduces the fundamentals of rarefied gas dynamics, including the Knudsen regimes of gas flow and the kinetic theory of dilute gases. Based on solving the Boltzmann equation, kinetic theory provides the basis for modelling non-equilibrium gas flows. The limitations of direct solution methods are discussed. The algorithm of the numerical DSMC method is then described, including the application of boundary conditions, the collision procedure, and the extraction of macroscopic properties. Some important features of the *dsmcFoamStrath* code (which is used to perform the DSMC simulations in this thesis) are also highlighted.

A detailed review of existing hybrid methodologies is provided in Chapter 3. After a brief discussion of 'unified' methods, focus is placed on continuum-DSMC hybrid methods. The capabilities and limitations of the four hybrid frameworks (DD, HMM, HMM-FWC, and IMM) are discussed. The coupling strategies adopted within these frameworks are examined in terms of their approach to spatial and temporal coupling, the information that is exchanged between the solvers, and the techniques for imposing boundary conditions on the micro sub-domains/elements. This review highlights the need for the continuum-DSMC method developed in this thesis.

Based on a HMM-FWC framework, the new continuum-DSMC hybrid method is presented in Chapter 4. The advantages of this method and possible applications are discussed. The multiscale methodology and the coupling strategy (i.e. the constraint of the micro elements and the correction of the continuum description) are described in detail. Finally, the iterative algorithm of the method is summarised.

In Chapter 5, the new hybrid method is validated for 1D energy coupling. This

is achieved by demonstrating its application to a micro Fourier flow problem for various rarefaction and temperature conditions. Results are compared with full DSMC simulations and the computational efficiency of the method is discussed.

Further validation is performed in Chapter 6, where the method is applied to high-speed micro Couette flow. With viscous heating, this flow problem requires simultaneous 1D momentum and energy coupling. Two wall speeds are considered to test the method's capabilities: one corresponding to supersonic flow, and the other to hypersonic flow. Once again, results are compared with full DSMC simulations.

In Chapter 7, the method is used to simulate the more realistic engineering problem of a rarefied gas flowing in a simple microscale crack. Essentially, the method is applied across the high-aspect-ratio cross-section of a microchannel that acts as a representative geometry; by assuming periodicity in the streamwise direction, the problem becomes 2D. Both isothermal and non-isothermal cases (i.e. where a temperature difference is present between the crack walls) are considered. The accuracy and computational cost of the method are again compared with equivalent full DSMC simulations.

Following a summary of the work presented in this thesis, Chapter 8 concludes with a discussion of possible future work.

1.4 Published papers

The work from this thesis has contributed to the following peer-reviewed papers:

1. Stephanie Y. Docherty, William D. Nicholls, Matthew K. Borg, Duncan A. Lockerby, and Jason M. Reese. Boundary conditions for molecular dynamics simulations of water transport through nanotubes, *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 228:186–195, 2014.

2. Stephanie Y. Docherty, Matthew K. Borg, Duncan A. Lockerby, and Jason M. Reese. Multiscale simulation of heat transfer in a rarefied gas, *International Journal of Heat and Fluid Flow*, 50:114–125, 2014.

3. Stephanie Y. Docherty, William D. Nicholls, Matthew K. Borg, and Jason M. Reese. Boundary conditions for molecular dynamics simulations of water transport through

carbon nanotubes, *9th International Conference on Heat Transfer, Fluid Mechanics and Thermodynamics (HEFAT)*, July 2012, Malta.

4. Stephanie Y. Docherty, Matthew K. Borg, Duncan A. Lockerby, and Jason M. Reese. Multiscale simulation of heat transfer in a rarefied gas, *5th International Conference on Heat Transfer and Fluid Flow In Microscale (HTFFM-V)*, April 2014, Marseille, France.

5. Stephanie Y. Docherty, Matthew K. Borg, Duncan A. Lockerby, and Jason M. Reese. A new heterogeneous multiscale technique for microscale gas flows, *4th Micro and Nano Flows (MNF) Conference*, September 2014, London, UK.

Chapter 2

Rarefied gas dynamics and the DSMC method

In this chapter, the fundamentals of rarefied gas dynamics will be discussed. This includes the classification of the different gas flow regimes, and the need for modelling techniques that are able to predict the flow behaviour of non-equilibrium gases. The kinetic theory of dilute gases will be introduced — as this is based on the solution of the Boltzmann equation, the limitations of direct solution methods will be considered. The importance of the numerical DSMC simulation method will then be explained, and its algorithm described in detail. Finally, the capabilities of the *dsmcFoamStrath* code (that will be used to perform all of the DSMC simulations in this thesis) will be reviewed.

2.1 Rarefied gas dynamics

2.1.1 Regimes of gas flow: the Knudsen number

Thermodynamic equilibrium in a gas is the result of highly frequent intermolecular and molecule-surface collisions. Low gas density, highly confined conditions, and large flow gradients drive a gas towards non-equilibrium. The level of non-equilibrium in a dilute gas is characterised by the Knudsen number,

$$Kn = \frac{\lambda}{l}, \tag{2.1}$$

where λ is the molecular mean free path of the gas (i.e. the average distance that gas molecules travel between successive collisions) and l is some characteristic length scale of the flow system. Setting l to a geometric dimension (e.g. a channel width or pipe diameter) results in a global Knudsen number Kn_{gl}, which provides an indication of the overall level of non-equilibrium in the system. Alternatively, if interest lies in a particular region of a system, a gradient-length local Knudsen number Kn_l can be defined by taking l as the scale of the gradient of a macroscopic quantity, e.g. $l = \rho/\nabla\rho$ where ρ is the gas density. Depending on Kn, gas flows fall into different regimes as indicated in Fig. 2.1.

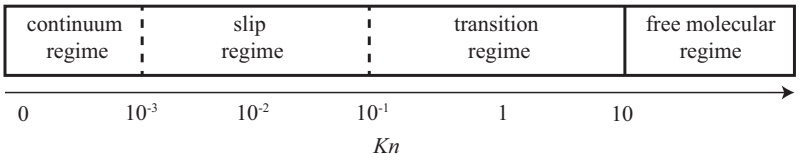

Figure 2.1: Gas flow regimes over different Knudsen numbers Kn. Note that there remains ambiguity in the upper and lower limits of the slip regime.

The continuum regime

In the limit $Kn \approx 0$, the frequency of molecular collisions is so high that the diffusive transport of heat and momentum can be neglected and the gas flow is considered to be in full thermodynamic equilibrium. From the continuum-fluid point of view, the flow is approximately isentropic and the conservation equations are the inviscid Euler equations.

In the range $0 < Kn < 10^{-3}$, the collision frequency is still very high and the flow is considered to be in quasi-equilibrium. The conventional Navier-Stokes-Fourier (NSF) equations (i.e. the conservation equations closed using the conventional linear NSF fluid-constitutive relations) are excellent for modelling the gas behaviour in this regime. The frequency of molecule-wall collisions is high enough that there are no discontinuities at gas-solid interfaces, and so the traditional 'no velocity slip' and 'no temperature jump' boundary conditions are valid.

The slip regime

As Kn increases, fewer molecular collisions occur. In the approximate range $10^{-3} < Kn < 0.1$, the molecule-wall collision frequency becomes so low that discontinuities appear — the macroscopic gas velocity and temperature at a solid surface are no longer equal to the velocity and temperature of the surface itself. The NSF model can, however, be extended into this regime by including Maxwell's velocity slip [33] and von Smoluchowski's temperature jump [34] boundary conditions.

The transition regime

When the Knudsen number rises above about 0.1, the collision frequency decreases further such that the gas is now far from thermodynamic equilibrium. Molecule-wall collisions become more frequent than intermolecular collisions, and momentum and/or thermal Knudsen layers form at the surfaces — these are essentially regions of non-equilibrium (with a thickness dependent on Kn) where the gas behaviour deviates from that of the linear NSF constitutive relations. Even with the use of slip and jump boundary conditions, the traditional NSF equations cannot model this phenomenon. The need for alternative modelling approaches in this flow regime has resulted in considerable research activity in recent years. Generally, these alternative approaches are based on the kinetic theory of dilute gases, which provides the foundation of the gas description in this flow regime; kinetic theory will be discussed further in section 2.1.2.

Note that this Kn range is a guide only. There is uncertainty in the literature regarding the validity of the linear NSF fluid-constitutive relations for Kn as large as 0.1 [35].

The free molecular regime

When $Kn > 10$, the flow becomes so rarefied that the molecules are separated by extremely large distances relative to their diameter. Intermolecular collisions are rare and molecule-wall collisions dominate. Analytical solutions to the Boltzmann equation can be obtained for simple geometries, as will be discussed in section 2.1.2. More complex geometries require numerical simulation (usually with MD) to obtain an accurate particle-surface interaction model for a given gas-surface combination.

2.1.2 Kinetic theory and the Boltzmann equation

Modelling techniques that are able to capture molecular behaviour are needed for gas flows that fall inside the transition regime. The most fundamental technique is deterministic MD, where the motion and interactions of individual molecules are governed by classical Newtonian mechanics. For a molecule i,

$$m_i \frac{d^2\mathbf{r}_i}{dt^2} = \sum_{\substack{j=1 \\ j\neq i}}^{N_{\text{mol}}} \mathbf{F}_{ij}, \tag{2.2}$$

where m_i is its molecular mass, $\mathbf{r}_i$ is its position, and the right hand side of this expression represents the net forcing on the molecule. The force $\mathbf{F}_{ij}$ acting between molecule i and another molecule j is a function of the intermolecular potential, most often based on a Lennard-Jones (LJ) model. As MD simulates all N_{mol} molecules in the system, the computational cost can be enormous; even for a dilute gas, a cubic centimetre volume at standard temperature and pressure contains 2.687×10^{19} molecules [36].

Although MD is physically accurate for any fluid, it is highly inefficient for modelling dilute gases. A gas is considered dilute when the mean molecular separation δ is considerably larger than the molecular diameter d, i.e. $\delta \gg d$. Intermolecular potentials can then be neglected and collisions are overwhelmingly likely to be binary. Rather than using a computationally costly deterministic approach, we can instead adopt a statistical technique based on the kinetic theory of dilute gases.

In kinetic theory we often refer to 'particles' — these can be atoms or molecules, or even groups of atoms/molecules. The overall aim is to estimate the probability of having a particle at a particular position, velocity, and internal state. Using both probability theory and statistical averaging, macroscopic property fields can be obtained from microscopic data. Note, however, that statistical fluctuations must be insignificant; the volume over which the property fields are averaged must contain a sufficient number of particles to reduce statistical error. The molecular chaos restriction ensures that the velocities of colliding particles are uncorrelated, requiring that the length scale for the averaging process is (at minimum) $100\times$ the average molecular separation [11]. Although kinetic theory is capable of describing polyatomic gases and gas mixtures, the

presence of internal degrees of freedom dramatically increases the model complexity. Such complexity is avoided in this thesis by considering simple monatomic dilute gases; each particle then has translational kinetic energy only.

Derived from the Liouville equation of statistical mechanics, the Boltzmann equation is the backbone of kinetic theory. Essentially, the equation describes the evolution of the (normalised) velocity distribution function f, which represents the distribution of particles as a function of position, velocity, and time. Assuming molecular chaos in a monatomic dilute gas, the Boltzmann equation can be expressed as [11],

$$\frac{\partial(nf)}{\partial t} + \mathbf{c}\frac{\partial(nf)}{\partial \mathbf{r}} + \mathbf{F}\frac{\partial(nf)}{\partial \mathbf{c}} = J(f, f^*), \tag{2.3}$$

where n is the number density, $\mathbf{F}$ is an external body force, and $\mathbf{r}$ and $\mathbf{c}$ are the position and velocity vectors of a particle, respectively. The term on the right hand side, $J(f, f^*)$, is the non-linear collision integral that accounts for the effects of particle collisions on the distribution function. With superscript * denoting post-collision values, this collision term can be written in full,

$$J(f, f^*) = \int_{-\infty}^{\infty}\int_{0}^{4\pi} n^2(f^* f_1^* - f f_1) c_r \xi d\Lambda d\mathbf{c}_1, \tag{2.4}$$

where f and f_1 are the two particle velocity distribution functions corresponding to $\mathbf{c}$ and $\mathbf{c}_1$, respectively, and c_r is the relative speed between the two colliding particles. The particle collision cross-section is denoted by ξ, while Λ is the solid angle (and $\xi d\Lambda$ is the differential cross-section). Upon solution of Eq. (2.3), weighted integrals of f can then be used to obtain macroscopic field properties, including the density, velocity, and temperature. Although the Boltzmann equation is perfectly valid for all Kn (provided that the assumptions of binary collisions and molecular chaos hold true), the non-linearity of the collision integral means that direct solution is very difficult.

In the continuum regime, the governing fluid conservation equations can, in fact, be derived from the Boltzmann equation by taking moments, i.e. by multiplying the Boltzmann equation in turn by the molecular mass, momentum, and energy, and then integrating over all velocity space (over all $\mathbf{c}$) [37]. In the limit $Kn \approx 0$, the gas is in

full equilibrium and the velocity distribution function takes the Maxwellian form f_o,

$$f_o(\mathbf{c}') = \left(\frac{m}{2\pi k_B T}\right)^{3/2} \exp\left[\frac{-m\mathbf{c}'^2}{2k_B T}\right], \tag{2.5}$$

where m is the molecular mass, k_B is Boltzmann's constant, and T is the gas temperature. In kinetic theory, a reference frame moving with the stream velocity is often adopted and so parameters are defined in terms of the thermal (or 'peculiar') velocities of the particles. The thermal velocity $\mathbf{c}'$ of a particle is its instantaneous velocity $\mathbf{c}$ relative to the stream velocity $\mathbf{c}_0$, i.e. $\mathbf{c}' = \mathbf{c} - \mathbf{c}_0$; the stream velocity $\mathbf{c}_0$ is equal to the mean of the instantaneous particle velocities. Note that we can express the Maxwellian speed distribution (which is visualised in Fig. 2.2 for a gas temperature of 273 K) in terms of the thermal speed c' [38], i.e.

$$f_o(c') = 4\pi \left(\frac{m}{2\pi k_B T}\right)^{3/2} c'^2 \exp\left[\frac{-mc'^2}{2k_B T}\right]. \tag{2.6}$$

The Maxwellian velocity distribution $f_o(\mathbf{c}')$ corresponds to a viscous stress tensor that is equal to the thermodynamic pressure and a heat flux that is zero, i.e. no diffusive momentum or energy transport; the conservation equations then reduce to the Euler equations, as mentioned previously.

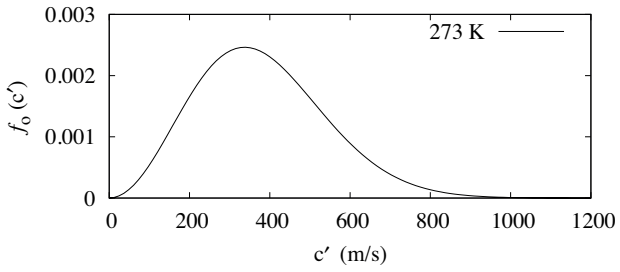

Figure 2.2: Maxwellian (thermal) speed distribution for a gas temperature of 273 K.

As we move away from thermodynamic equilibrium, the particle distribution diverges from Maxwellian form. If the departure from equilibrium is only slight, the Chapman-Enskog (CE) technique can be employed — this aims to obtain an analytical solution to the Boltzmann equation by considering a distribution function which is a small perturbation from the Maxwellian form. Using a Kn series expansion, this

distribution can be expressed as,

$$f = f_o\{1 + b_1 Kn + b_2 Kn^2 + ... + b_n Kn^n\}, \tag{2.7}$$

where coefficients $b_{1,2,...n}$ are functions of density, velocity and temperature only [39]. Note that the zeroth-order approximation in Kn results in the Maxwellian distribution, and so yields the Euler equations. The first-order approximation in Kn corresponds to the linear constitutive relations of the NSF conservation equations. The second-order approximation produces the complex non-linear constitutive relations of the Burnett equations [40], while the third-order approximation results in the super-Burnett equations [41]. As the Burnett and super-Burnett equations are accurate to a higher order in Kn, they could essentially extend the continuum-fluid description beyond the NSF equations and into the transition regime. However, the departure from equilibrium must be small for the CE to be valid, and the mathematical complexity involved in the solution of these equation sets means that applications are restricted to the most simple geometries. In addition, there are some concerns regarding the stability and thermo-mechanical consistency of these higher order Burnett approaches [42].

Grad [43] took a somewhat different approach, expanding the distribution function as a series of Hermite polynomials, where the coefficients are combinations of the moments of f. With infinite Hermite coefficients, no molecular information is lost. However, for practicality, the distribution function must be truncated: Grad truncated at the third order in Hermite polynomials, resulting in his well-known 13 moment equations (G13). Based on Grad's work, there has been significant development of moment methods in recent years. In order to smooth out artificial discontinuities of the G13 equations, Struchtrup & Torrilhon [44] derived a set of regularized 13 moment equations (R13) by using an additional Chapman-Enskog-like expansion. Gu & Emerson [45] truncated the distribution function at the fourth order in Hermite polynomials and derived a set of 26 moment equations that were regularized by the procedure of Struchtrup & Torrilhon. Despite these developments, a number of problems remain, including the presence of instabilities and the need for additional boundary conditions [39]. Although recent studies have proposed new boundary conditions [46–48] and have considered implementation for problems like Couette and Poiseuille flows [49–52],

further development of moment methods is needed for complex 3D flows.

Generally, analytical solution of the Boltzmann equation is restricted to free molecular flows around simple geometries; analytical solution is not possible for arbitrary complex geometries in the transition regime. Attention has therefore turned to numerical solution methods. Unfortunately, with the high dimensionality of the collision term, direct numerical solution is extremely computationally expensive (particularly for 3D systems) due to the need for discretisation of both physical space and velocity space. Simplification of the collision integral makes numerical solution much more accessible. The Bhatnagar-Gross-Krook (BGK) method [53], for example, replaces the collision integral in Eq. (2.4) with a simplified representation of the collision processes — this effectively linearises the Boltzmann equation. Techniques such as the discrete velocity method (DVM) can then be implemented for solution [54–56]. However, simplifying the collision term also reduces the applicability of such solution methods, with studies focusing mainly on simple benchmark flow problems. The standard BGK model, for example, is rarely accurate for complex, highly non-equilibrium flows [35], and modifications are required in order to obtain the proper Prandtl number [57].

In recent years, the rise in parallel computing technologies has lead to a renewed focus on numerical solution methods. In particular, there has been an interest in using Graphic Processing Units (GPUs) to perform faster calculations [58, 59]. Despite this, numerical Boltzmann solvers are not yet regarded as a viable alternative to particle-based simulation methods. Today, the most popular numerical tool in rarefied gas dynamics is the particle-based direct simulation Monte Carlo (DSMC) method.

2.2 The DSMC method

The DSMC method was first proposed by Bird [60] in the 1960s as an alternative to expensive MD simulations. Although it is based on kinetic theory (and so is subject to the same dilute gas and molecular chaos restrictions), DSMC does not aim to solve the Boltzmann equation directly. Instead, the aim is to simulate the physics of the system. The fundamental concept is to track a large number of particles that move through a computational mesh, storing their position, velocity, and internal state. During a simulation, the particles collide with each other and with bounding surfaces while con-

serving mass, momentum, and energy. The main advantage of the DSMC method over MD is the computational savings that are offered by two key assumptions: first, each simulated particle is able to represent a large number of real gas molecules; secondly, particle collisions can be treated stochastically rather than deterministically. Essentially, DSMC decouples the particle motion and collisions over small time intervals — particle movements are computed deterministically depending on the velocity vector and DSMC time step, after which collisions are handled statistically.

The DSMC method became very popular in the 1970s due to a focus on the development of hypersonic aerodynamics. This led to rapid advances in the method, including the extension to gas mixtures [61] and diatomic molecules [62]; the introduction of the Larsen-Borgnakke model [63] for redistribution of energy between translational and rotational modes during inelastic collision; Bird's No Time Counter (NTC) scheme [64] for obtaining accurate collision rates; and the implementation of the sub-cell approach [11] to promote nearest neighbour collisions. For several decades, there was concern over the relationship between the method and the Boltzmann equation [65]. This was, however, put to rest in 1992 when it was proved mathematically that DSMC can produce a solution to the Boltzmann equation for a monatomic gas in the limit of an infinite number of particles [66]. It is now widely accepted that DSMC is the most robust method for modelling highly rarefied dilute gas flows in complex geometries. Results from newer modelling approaches (e.g. analytical methods and hybrid approaches) are often validated against results from full DSMC simulations.

2.2.1 The DSMC algorithm

Before beginning a DSMC simulation, a computational mesh is generated over the flowfield — this will be needed for selecting particle collision partners. Cell dimensions are set to a fraction (usually 1/3) of the mean free path to ensure realistic collisions. The boundary conditions that are to be imposed at the outer limits of the system are then defined. The molecular properties of the gas are also specified, including the molecular mass m and the number of real molecules F_N that each DSMC particle will represent. The molecular collision model is selected; the popular Variable Hard Sphere (VHS) model is used for all DSMC simulations in this thesis. This model, unlike the

fixed-diameter Hard Sphere (HS) model, is able to preserve the viscosity-temperature relationship of a gas by allowing a variable molecular diameter d_{VHS} [11], i.e.

$$d_{\mathrm{VHS}} = \left(\frac{15m(RT/\pi)^{1/2}}{2\mu(7-2\omega)(5-2\omega)} \right)^{1/2}, \tag{2.8}$$

where μ is the gas dynamic viscosity, and ω is the viscosity exponent of the gas. The VHS mean free path is then calculated from,

$$\lambda = \frac{1}{\sqrt{2}\pi d_{\mathrm{VHS}}^2 n}, \tag{2.9}$$

where n is the gas number density.

After defining these simulation parameters, the standard DSMC algorithm can begin; this algorithm is summarised in Fig. 2.3. First, the system is initialised: the mesh is populated with DSMC particles at an equilibrium state, i.e. the thermal velocity of each particle is sampled from a Maxwellian distribution at the desired macroscopic temperature and velocity. The main loop of the algorithm then consists of four steps.

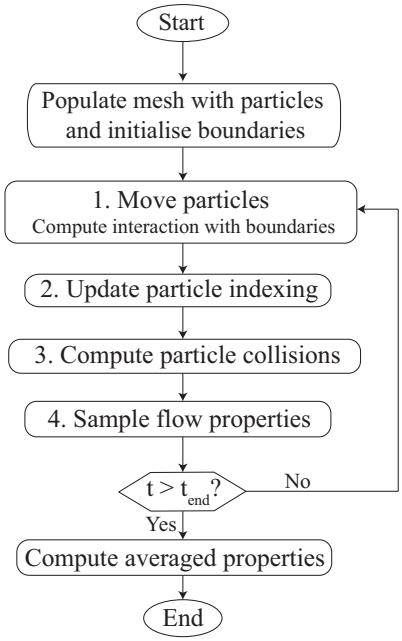

Figure 2.3: Flowchart of the standard DSMC algorithm.

1. Particle movements

The first step is to move each particle i along its velocity vector $\mathbf{c}_i$ for a duration of one DSMC time step δt, assuming that the velocity remains constant and that there are no collisions. If the decoupling of movements and collisions is to produce an accurate solution, δt must be significantly smaller than the average time between gas particle collisions, referred to as the mean collision time t_{mc} [11].

Particle-boundary interactions are also dealt with in this step. There are three important types of boundary condition that are commonly imposed in DSMC simulations. The first of these deals with the flow of particles across the boundaries (into or out of the system) to obtain the desired freestream conditions. Inflow boundaries are based on the insertion of particles depending on the equilibrium Maxwellian number flux [11], while particles that make contact with outflow boundaries are instantly deleted from the system. Pressure-based inflow/outflow boundary conditions have also been developed especially for low-speed flows through microgeometries [67].

The second type of boundary condition deals with particle-surface interactions. In DSMC, the most frequently used interaction model is Maxwell's [33], which considers both specular and diffuse reflections as illustrated in Fig. 2.4. Specular reflection is representative of a perfectly smooth surface with perfect slip: incident particles conserve tangential momentum — the normal component of their incident velocity is reversed while their tangential velocity and energy remain unchanged, and the angle of reflection Θ_{refl} is the same as the angle of incidence Θ_{inc}. On the other hand, diffuse reflection is representative of a microscopically rough surface, i.e. incident particles lose all of their tangential momentum to the wall. The post-interaction velocity does not depend on the incident velocity; it is sampled from a Maxwellian distribution depending on the surface temperature, with equal probability in all directions. For a given gas-boundary interaction, the fraction of particles that experience a diffuse reflection α is known as the tangential momentum accommodation coefficient (TMAC), and depends on the gas properties, the surface material, and the surface roughness. While Maxwell's model is sufficient for the simple test cases of this thesis, future cases may benefit from the use of the more sophisticated Cercignani-Lampis-Lord (CLL) interaction model [68]. However, there remains uncertainty in the type of problems that require the CLL model.

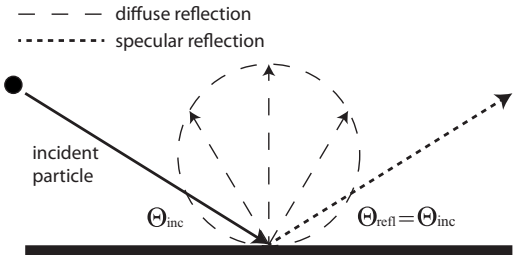

Figure 2.4: Schematic of Maxwell's diffuse and specular particle-surface interactions..

The third boundary type is periodic — when a particle crosses a periodic boundary, it simply re-enters at the mirror boundary with the same velocity and internal state.

2. Indexing

After moving the particles and computing boundary interactions, each particle is indexed into a cell of the computational mesh depending on its current position. This indexing is required as particle collision partners are selected on a cell-by-cell basis.

3. Particle Collisions

Intermolecular collisions are computed probabilistically. It is vital that the collision process is representative of real gas interactions and so particles should collide only with neighbouring particles. For this reason, particle collision pairs are selected from the same cell. Nearest neighbour collision can be further encouraged by using a sub-cell technique (or a more advanced transient-adaptive sub-cell technique [69]), in which particles are restricted to collisions with other particles that reside in the same sub-cell.

To simulate realistic collision rates, an appropriate collision modelling scheme must be chosen. Bird's NTC scheme can produce accurate collision rates in exchange for minimal computational effort, and so has become the standard approach in DSMC. In this scheme, a number of particle pairs are selected at random within each cell or sub-cell, regardless of their position. An acceptance-rejection approach is then adopted. For each particle pair, a collision is accepted if

$$\frac{(\xi_{\mathrm{tot}} c_r)}{(\xi_{\mathrm{tot}} c_r)_{\mathrm{max}}} > R_f, \tag{2.10}$$

where ξ_{tot} is the total collision cross-section (i.e. $\xi_{\text{tot}} = \pi d^2$), c_r is the relative speed between the pair, and R_f is a random fraction between 0 and 1. Initially, $(\xi_{\text{tot}} c_r)_{\max} = (\xi_{\text{tot}} c_r)$; subsequently, if $(\xi_{\text{tot}} c_r)$ is found to be greater than $(\xi_{\text{tot}} c_r)_{\max}$, then the maximum value is updated. After the collision is processed, or if the pair are rejected, a new particle pair is selected and the acceptance-rejection procedure is repeated. This continues until the required number of candidate pairs,

$$M_{\text{cand}} = \frac{N_c(N_c - 1)F_N(\xi_{\text{tot}} c_r)_{\max}\delta t}{2V_c}, \tag{2.11}$$

within the cell volume V_c at the given time step δt has been selected (where N_c is the instantaneous number of DSMC particles in the cell).

In DSMC, accepted collisions are processed by altering the velocities (but not the positions) of both particles. For the monatomic gases considered in this thesis, particles i and j will exchange translational energy only and so the collision will be fully elastic. Conserving momentum, the post-collision velocities for each particle can be calculated from,

$$\begin{aligned} \mathbf{c}_i^* &= \mathbf{c}_m^* + \left(\frac{m_j}{m_i + m_j}\right)\mathbf{c}_r^*, \\ \mathbf{c}_j^* &= \mathbf{c}_m^* - \left(\frac{m_i}{m_i + m_j}\right)\mathbf{c}_r^*, \end{aligned} \tag{2.12}$$

where superscript * denotes post-collision values. The centre of mass velocity $\mathbf{c}_m$ remains unchanged by the collision, i.e.

$$\mathbf{c}_m^* = \mathbf{c}_m = \frac{m_i\mathbf{c}_i + m_j\mathbf{c}_j}{m_i + m_j}. \tag{2.13}$$

The post-collision relative velocity $\mathbf{c}_r^*$ is computed from,

$$\mathbf{c}_r^* = c_r^*[(\sin\vartheta\cos\varphi)\hat{\mathbf{x}} + (\sin\vartheta\sin\varphi)\hat{\mathbf{y}} + (\cos\vartheta)\hat{\mathbf{z}}]. \tag{2.14}$$

From the conservation of energy, it can be deduced that the magnitude of the relative velocity between the particles c_r is not changed by the collision, i.e. $c_r^* = c_r = \|\mathbf{c}_i - \mathbf{c}_j\|$. With the isotropic scattering law of the VHS model, the scattering angles ϑ and φ are

uniformly distributed over a unit sphere, between 0 and 2π radians. The azimuthal angle φ is equal to $2\pi R_f$, while the elevation angle $\vartheta = \cos^{-1}(2R_f - 1)$.

For diatomic or polyatomic gases, each particle pair should be considered for rotational energy exchange and collisions may be inelastic; the phenomenological Larsen-Borgnakke model can then be used to compute post-collision velocities.

4. Sampling

When modelling real applications, engineers are not generally interested in microscopic information — interest lies in the variation of the macroscopic properties. Macroscopic information is recovered in DSMC from statistical averaging of the microscopic information, using expressions taken from kinetic theory [11]. We restrict the discussion here to a monatomic simple gas, where all particles are of the same chemical species and there are no effects of rotational or vibrational energy. The discussion would be more complex for a gas mixture with multiple chemical species.

First, we discuss the extraction of macroscopic properties by averaging over all particles N_V in a volume V. Typically, this volume is a DSMC computational cell, however the same process applies to any user-defined volume. The gas number density n can be computed from,

$$n = \frac{N_V F_N}{V}. \tag{2.15}$$

Note that it is important to account for the fact that each DSMC particle represents some number F_N of real gas molecules. The mass density ρ is then simply,

$$\rho = nm, \tag{2.16}$$

where m is the molecular mass. As mentioned in section 2.1.2, the macroscopic stream velocity $\mathbf{c}_0$ (also denoted $\mathbf{u}$ in this thesis) is the mean of the instantaneous particle velocities $\mathbf{c}$, i.e.

$$\mathbf{u} = \mathbf{c}_0 = \overline{\mathbf{c}} = \frac{1}{N_V}\sum_{i=1}^{N_V} \mathbf{c}. \tag{2.17}$$

The use of an overline indicates the average value of the expression over all particles in the sample. If a reference frame moving with this stream velocity is adopted, the

macroscopic properties are then defined in terms of the particle thermal velocities $\mathbf{c}'$, i.e. $\mathbf{c}' = \mathbf{c} - \mathbf{c}_0$. The pressure tensor $\mathbf{p}$ is computed from,

$$\mathbf{p} = \rho\overline{\mathbf{c}'\mathbf{c}'}, \tag{2.18}$$

for example, in component form,

$$p_{xy} = \rho\overline{u'v'}. \tag{2.19}$$

The macroscopic scalar pressure p is then the average of the three normal components of the pressure tensor, i.e.

$$p = \frac{1}{3}\rho\left(\overline{u'^2} + \overline{v'^2} + \overline{w'^2}\right). \tag{2.20}$$

The components of the viscous stress tensor $\boldsymbol{\tau}$ are computed from the negative of the pressure tensor, with the scalar pressure subtracted from the normal components, e.g.

$$\tau_{xx} = -\left(\rho\overline{u'^2} - p\right), \tag{2.21}$$

and,

$$\tau_{xy} = -\left(\rho\overline{u'v'}\right). \tag{2.22}$$

In monatomic gases, the particles have only translational kinetic energy ε_{tr}. Consequently, the overall gas temperature T is equal to its translational temperature T_{tr},

$$T_{tr} = \frac{2}{3k_B}\varepsilon_{tr} = \frac{1}{3k_B}m\left(\overline{u'^2} + \overline{v'^2} + \overline{w'^2}\right). \tag{2.23}$$

With no rotational or vibrational energy, the components of the heat flux vector $\mathbf{q}$ can be computed from, for example,

$$q_x = \frac{1}{2}\rho\overline{c'^2u'}. \tag{2.24}$$

The macroscopic properties of the gas at the bounding walls will also be needed in the later chapters of this thesis. These can be obtained by summing over all particles

N_{sur} that strike the surface [70]. For example, for a horizontal wall, the speed of the gas u_{gs} at the surface (i.e. the tangential component of the stream velocity at the surface, $c_{0_{\text{gs},t}}$) is computed from,

$$u_{\text{gs}} = c_{0_{\text{gs},t}} = \frac{\sum_{j=1}^{N_{\text{sur}}} \left[(m/|c_{n,j}|)\, c_{t,j}\right]_{\text{tot}}}{\sum_{j=1}^{N_{\text{sur}}} \left[(m/|c_{n,j}|)\right]_{\text{tot}}}, \tag{2.25}$$

where $c_{n,j}$ and $c_{t,j}$ are the normal and tangential velocity components of particle j, respectively. The subscript 'tot' signifies that the summation is performed both before and after the gas-surface interaction. The temperature of the gas at the wall surface T_{gs} is then obtained from,

$$T_{\text{gs}} = \frac{1}{3k_B} \frac{\sum_{j=1}^{N_{\text{sur}}} \left[(m/|c_{n,j}|)\left(\|\mathbf{c}_j\|^2\right)\right]_{\text{tot}} - \sum_{j=1}^{N_{\text{sur}}} \left[(m/|c_{n,j}|)\right]_{\text{tot}} u_{\text{gs}}^2}{\sum_{j=1}^{N_{\text{sur}}} \left[(1/|c_{n,j}|)\right]_{\text{tot}}}, \tag{2.26}$$

where $\|\mathbf{c}_j\|$ is the magnitude of the particle velocity. At each time step δt, the tangential shear stress components of the gas at the wall $\tau_{\text{gs},t}$ can be sampled from the difference in the tangential momentum flux before (denoted by subscript 'inc') and after (denoted by subscript 'refl') the gas-surface interaction [38], i.e.

$$\tau_{\text{gs},t} = \frac{F_N}{A\delta t} \sum_{j=1}^{N_{\text{sur}}} \left[(mc_{t,j})_{\text{refl}} - (mc_{t,j})_{\text{inc}}\right], \tag{2.27}$$

where A is the surface sampling area. Similarly, the heat flux on the surface (i.e. the normal component of the heat flux vector at the wall, $q_{\text{gs},n}$) is sampled from the difference in the energy fluxes before and after the gas-surface interaction at each time step, i.e.

$$q_{\text{gs},n} = \frac{F_N}{A\delta t} \sum_{j=1}^{N_{\text{sur}}} \left[\left(\frac{1}{2}mc_j^2\right)_{\text{inc}} - \left(\frac{1}{2}mc_j^2\right)_{\text{refl}}\right]. \tag{2.28}$$

Unfortunately, statistical scatter is inherent in DSMC. For steady-state problems like those considered in this thesis, this problem can be lessened by using cumulative time averaging. When a simulation reaches steady-state, the gas properties are averaged

over a large number of time steps $g = 1, 2, ...G_{\text{av}}$, for example,

$$\overline{n} = \frac{1}{G_{\text{av}}} \sum_{g=1}^{G_{\text{av}}} n(g). \tag{2.29}$$

For transient problems, scatter reduction can be achieved from ensemble averaging (i.e. repeating the simulation many times).

2.2.2 *dsmcFoamStrath*

In collaboration with researchers at the University of Strathclyde, OpenCFD Ltd have released MD and DSMC solvers within the framework of their C++ toolbox, OpenFOAM [71]. Primarily a Computational Fluid Dynamics (CFD) package, OpenFOAM is completely open source under the GNU general public licence and is freely available for download. The DSMC solver, named *dsmcFoam*, includes the key features of any modern DSMC code: particle initialisation in arbitrary 2D/3D geometries; particle tracking in structured and unstructured meshes; freestream and periodic boundaries; Maxwellian diffuse/specular gas-surface interactions; automatic sub-celling; NTC partner selection; Larsen-Borgnakke rotational energy redistribution; steady-state and transient simulation capabilities; and unlimited parallel processing capability.

The open source nature of OpenFOAM means that users are also able to extend the code to suit their own needs. The standard *dsmcFoam* code has been extended within the research group of the author [36] to provide additional functionality; additions include a vibrational mode, quantum-kinetic (QK) chemical reactions, and a CLL gas-surface interaction model. This extended solver, referred to as *'dsmcFoamStrath'*, has been validated for various benchmark cases, including hypersonic and microchannel flows [72–74]. It is therefore used to perform all of the hybrid and full-scale DSMC simulations in this thesis.

Measurement techniques

As discussed in the previous section, the macroscopic properties in a volume can be extracted by averaging over all particles in that volume according to Eqs. (2.15) – (2.24). The standard *dsmcFoam* distribution simply averages on a cell-by-cell basis.

Driven by the needs of the hybrid method presented in the later chapters of this thesis, a new measurement framework [36] has been added to the *dsmcFoamStrath* code; this framework includes two alternative measurement techniques — the 'method of zones' (MOZ) and the 'method of bins' (MOB) are based on averaging over the particles in some user-defined volume.

In the MOZ technique, a user-defined 'zone' is created. Consisting of connected DSMC computational cells, this zone can be any size or shape and can be at any location on the DSMC mesh. The spatial average of the macroscopic properties in that zone is then extracted; this technique cannot, however, capture spatial gradients.

Spatial gradients can be captured by the MOB technique. A volume (the full domain or a user-defined zone) is split into a number of 'bins' and the spatial average of the properties in each bin is extracted. In this thesis, these measurement bins are 1D or 2D, both shown in Fig. 2.5. Other possibilities include 2D radial bins (for cylindrical geometries) or even 3D bins. It is important to note that this binning method is independent of the computational mesh used by the DSMC algorithm — this provides the ability to define the resolution of the extracted macroscopic fields, and to control the level of noise, without affecting the particle collision rate.

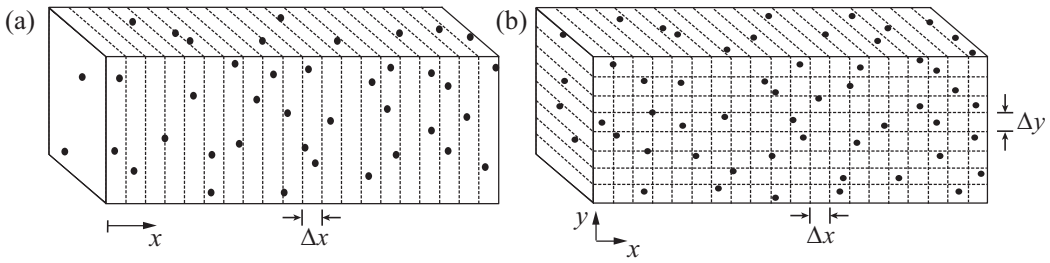

Figure 2.5: Schematic of the MOB using (a) 1D planar bins and (b) 2D bins.

Note that it would be much too memory-intensive to store particle properties at each time step and then calculate the macroscopic properties on completion. Instead we perform measurements 'on the fly', storing less memory-intensive macroscopic properties at each write interval. All of the measurement tools in *dsmcFoamStrath* can perform instantaneous measurements until steady-state is reached (resetting after write out), and cumulative measurements during the steady-state run (averaging over time).

State controllers

Generally, DSMC simulations are constrained by imposing conditions at the system boundaries. In MD, constraints can also be applied by implementing state controllers [75, 76] in some prescribed volume; a desired macroscopic state is obtained by controlling the molecular behaviour. Again motivated by the requirements of the hybrid method of later chapters, this methodology has been extended to the DSMC method. State controllers for the density, velocity, and temperature have been added to the *dsmcFoamStrath* code with the help of others from the research group of the author.

A simple feedback loop algorithm is used for all three controllers: in the volume of interest the measured property is compared against its target value, with the difference providing the basis for the controlling action. By calculating the difference between the target density ρ_{tar} and the density measured ρ_{meas} in the volume V, the density controller is able to insert or delete particles accordingly, i.e.

$$\Delta N = (\rho_{\text{tar}} - \rho_{\text{meas}}) \frac{V}{F_N m}, \tag{2.30}$$

where ΔN is the number of particles to be inserted/deleted (note that this is rounded to the nearest integer). Inserted particles are given an arbitrary position within the volume, with their thermal velocities sampled from a Maxwellian distribution at the desired macroscopic conditions. If particles must be deleted from the volume, these are selected at random. Velocity control is fairly straightforward; the difference between the target velocity $\mathbf{u}_{\text{tar}}$ and the measured stream velocity $\mathbf{u}_{\text{meas}}$ (i.e. $\mathbf{c}_{0,\text{meas}}$) is added to the original velocity $\mathbf{c}$ of each particle in the volume, i.e.

$$\mathbf{c}_{\text{new}} = \mathbf{c} + (\mathbf{u}_{\text{tar}} - \mathbf{u}_{\text{meas}}). \tag{2.31}$$

The temperature controller is based on a Berendsen thermostat [77], which rescales the thermal velocities of the particles in the volume according to a scaling factor, i.e.

$$\mathbf{c}_{\text{new}} = (\mathbf{c} - \mathbf{u}_{\text{meas}}) \left[1 + \frac{\delta t}{\varsigma} \left(\frac{T_{\text{tar}}}{T_{\text{meas}}} - 1 \right) \right]^{1/2} + \mathbf{u}_{\text{meas}}, \tag{2.32}$$

where T_{tar} is the target temperature, T_{meas} is the temperature measured in the volume,

and ς is a relaxation constant that determines the speed of the response.

A flexible time scheme is used, meaning that the frequencies of measurement and control are independent of each other; each can be specified to suit the given state property and the particular problem. It may be useful, for example, to measure a property less frequently than it is controlled in order to reduce the statistical inaccuracies fed back to the controller algorithm. Note that if a system requires simultaneous density, velocity and temperature control, then these are implemented sequentially. Density control is performed first, followed by velocity control. Temperature control is then performed on the updated particle velocities from the velocity control.

Based on the measurement framework of the *dsmcFoamStrath* code, the state controllers can be implemented using the MOZ or MOB techniques. The MOZ enables control of the spatial average of a property in a user-defined zone of any shape/size. With the MOB, spatial gradients can be imposed across a user-defined zone — the volume is divided into a number of bins and the spatial average of each bin is then controlled. Like the measurement bins, the control bins in this thesis are either 1D planar bins or 2D bins (see Fig. 2.5). As the MOB is completely independent of the computational mesh used by DSMC, the resolution of the imposed macroscopic fields can be defined without affecting the DSMC computations.

2.3 Summary

There is a need for techniques that are able to accurately model the flow behaviour of dilute gases in the transition regime. At present, direct (analytical or numerical) solution of the Boltzmann equation is infeasible for complex 3D flow problems and/or where the departure from equilibrium is significant. The numerical DSMC method is currently the most effective (and most popular) simulation tool for non-equilibrium dilute gas flows — it is considerably more efficient than MD for these types of flows, and the macroscopic flow properties can be easily obtained from the microscopic information. The *dsmcFoamStrath* code that has been developed within the framework of the OpenFOAM toolbox is a robust DSMC solver; as well as having the important features of modern DSMC codes, this code now includes measurement and control capabilities that will be useful in subsequent chapters of this thesis.

Chapter 3

Hybrid methods for dilute gases

The DSMC method is valid for dilute gases across all Knudsen regimes, from the continuum to the free-molecular. However, despite it being significantly more efficient than MD, the computational cost of tracking hundreds of thousands (or even millions) of DSMC particles, computing collisions between these particles, and performing sampling of the property fields is still enormous. This cost is raised further by the need to reduce statistical scatter in the sampled fields, either by cumulative time averaging or ensemble averaging. Even with advances in high performance computing and parallel processing, realistic engineering problems can require months or even years of simulation time.

Fortunately, the great majority of conventional macroscale flow problems can be described using a continuum-fluid model, and so can be resolved through the traditional discretisation techniques of CFD. In some cases, however, there may be regions of the flow where the molecular behaviour at the microscale impacts flow behaviour at the macroscale. This can occur, for example, when there are large gradients in fluid properties, or when surface effects become dominant. While traditional CFD cannot capture the important microscale phenomena, it would be computationally prohibitive to implement the DSMC method across the entire flowfield. This has led to proposals for multiscale 'hybrid' methods. For dilute gases, hybrid methods can be split into two categories: unified methods, and continuum-DSMC methods.

3.1 Unified methods

The basis of a unified approach is that the flowfield is divided into equilibrium and non-equilibrium regions, and the same type of description (kinetic or particle-based) is used in both — the key advantage of this is that the transfer of information between the regions is naturally accommodated. An example of an 'all-kinetic' unified solver is that of Kolobov *et al.* [78, 79], in which a numerical Boltzmann solver is applied in non-equilibrium regions and the equilibrium regions are treated using a kinetic CFD scheme. Another example of an all-kinetic solver is the unified gas-kinetic scheme of Xu and Huang [80], which is based on the use of the BGK model everywhere.

'All-particle' unified hybrids have also been developed, for example by Tiwari and Klar [81]. In these schemes, a DSMC solver is implemented in the non-equilibrium regions, while the particle approach that is applied in the equilibrium regions avoids the DSMC collision computation, instead redistributing particle velocities under the assumption of local equilibrium. Unfortunately, this type of equilibrium particle approach is often plagued by numerical diffusion errors due to unrealistic random particle motion. Burt and Boyd [82, 83] were, however, able to reduce numerical diffusion errors and statistical scatter with a low-diffusion equilibrium particle method: by ensuring that the particles move along trajectories similar to the gas streamlines, the random thermal motion was suppressed.

Although simplifications to the Boltzmann equation in the equilibrium regions mean that all-kinetic hybrids are more efficient than applying a typical numerical Boltzmann solver everywhere, using a typical Boltzmann solver in the non-equilibrium regions is significantly more difficult (and possibly more expensive) than using a DSMC solver, particularly for complex 3D problems. Also, despite all-particle hybrid methods being less computationally demanding than a full DSMC simulation, the use of a particle method in regions of equilibrium flow is much less efficient than a conventional continuum-fluid solver. Generally, a traditional continuum-fluid solver is the most efficient modelling approach for regions of equilibrium flow and, at present, the DSMC method is the most accessible method for modelling non-equilibrium flows. Continuum-DSMC hybrid methods have therefore been of increased interest in recent years.

3.2 Continuum-DSMC hybrid methods

As well as combining the efficiency of a continuum approach with the accuracy of the DSMC method, continuum-DSMC methods have the additional advantage that both types of solver are fairly mature and so have been heavily validated and optimised. However, difficulties exist in coupling two solvers that are so different in nature. In 2004, Wijesinghe and Hadjiconstantinou [17] highlighted two significant challenges in the development of this type of method.

The first challenge is the selection of the coupling technique. The most efficient method will exploit spatial and temporal scale separation as much as possible; as mentioned in section 1.1, scale separation exists where the microscopic behaviour is only loosely coupled with the behaviour on a larger scale. Note that the efficiency of a hybrid method can be measured by the computational speed-up offered in comparison with a full particle simulation of the given flow problem. Hybrid methods in the literature have varying abilities to exploit scale separation depending on: the type of coupling framework being used; the spatial configuration of the solvers; the type of information being exchanged; and the frequency of the coupling exchange. Selection of the appropriate coupling technique is therefore a complex issue.

The second challenge cited in [17] lies in the application of boundary conditions at the interface between the solvers. It is particularly problematic to translate the macroscopic continuum variables into boundary conditions that are suitable for a particle simulation. Proposed approaches are often based on the use of a particle reservoir, where the local continuum property fields are used to impose a Maxwellian or Chapman-Enskog (CE) particle distribution. As macroscopic fields can be extracted fairly easily from a DSMC solver by statistical averaging, imposing boundary conditions on the continuum domain is much less problematic. The level of statistical scatter in the particle information must, however, be low to avoid any coupling instabilities.

To date, there is no one continuum-DSMC hybrid method that can be universally applied — the methods proposed in the literature have often been designed to deal with particular types of flow problem. Generally, two popular hybrid frameworks have emerged for both continuum-MD (used for dense fluid flows) and continuum-DSMC coupling: the domain decomposition (DD) technique, and the Heterogeneous Multiscale

Method (HMM). Derived from the HMM, the HMM with field-wise coupling (HMM-FWC), and the Internal-flow Multiscale Method (IMM) have also been proposed in recent years. In the following sub-sections, the capabilities and limitations of these four frameworks will be discussed. Focusing mainly on continuum-DSMC methods, the coupling strategies adopted within these frameworks will also be discussed, as well as the techniques used for imposing boundary conditions on the particle sub-domains. This review of existing hybrid methods will provide a backdrop for the HMM-FWC continuum-DSMC method that will be presented in the following chapters of this thesis, highlighting its benefits and capabilities as well as possible future developments.

3.2.1 Domain decomposition (DD)

The DD framework exploits spatial scale separation by dividing the flowfield into separate sub-domains: a full continuum-fluid model (closed by conventional constitutive relations) is used to resolve the macroscale phenomena in a large 'macro' sub-domain, while a particle/molecular solver resolves the microscale phenomena in a small 'micro' sub-domain. These micro and macro sub-domains are independent but communicate through some coupling interface or overlap. Generally, the DD approach is very popular as it is able to take advantage of the fact that non-equilibrium flow most often occurs close to bounding surfaces; a micro sub-domain can be placed in this region of non-equilibrium while the remainder of the system forms the macro sub-domain.

Based on the type of information that is exchanged between the solvers, DD continuum-DSMC techniques in the literature can be split into two categories: flux-based coupling and state-based coupling. A schematic of both is shown in Fig. 3.1. Typical flux-based coupling methods match fluxes of mass, momentum, and energy at the coupling interface. The fluxes in the DSMC solver are computed from the particles crossing the interface, while the fluxes in the continuum solver must be extrapolated from the macroscopic state properties and their gradients. These fluxes are then modified to ensure conservation across the interface and are applied as boundary conditions, updating the continuum solver and providing the basis for generating the chosen distribution function in the DSMC solver [84].

In state-based coupling, an overlap is used to match the macroscopic state quan-

tities, like velocity and temperature; each solver essentially provides the other with a Dirichlet boundary condition. Averaged state quantities extracted from the DSMC cells of the overlapping region are passed to the 'continuum boundary cells' to update the continuum solution. Simultaneously, the average state quantities extracted from the continuum cells of the overlapping region are passed to the 'DSMC boundary cells', where they are used to generate the chosen particle distribution.

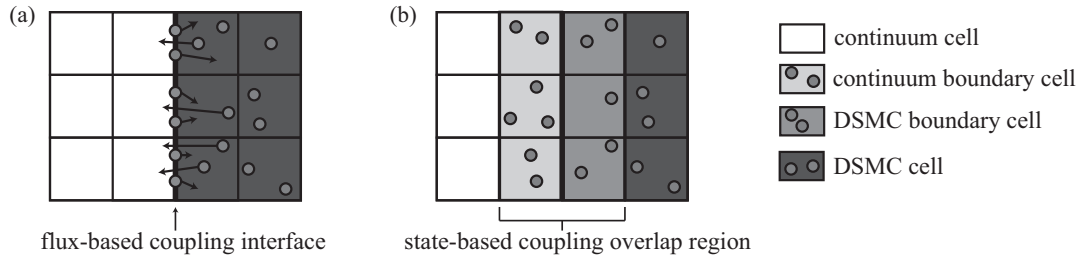

Figure 3.1: Schematic of typical (a) flux-based and (b) state-based coupling techniques for the DD framework. Note that the continuum-fluid and DSMC cells are shown to be of the same dimensions but this is not necessary.

Flux-based methods lend themselves to unsteady flows [17]. The fluxes at the interface are typically exchanged every time step, with both the continuum-fluid solver and the particle solver using the same time step [21]; this is often referred to as 'strong', 'full', or 'tight' temporal coupling, and is shown schematically in Fig. 3.2(a). However, a major issue with flux-based coupling is its sensitivity to the high level of statistical scatter associated with the measurement of fluxes in particle-based simulation methods. This was highlighted in 1992 by Wadsworth and Erwin [12] in the application of their strongly coupled, flux-based NS-DSMC method to 2D unsteady slit flow — while the 'continuum-to-particle' flux was time-explicit, cumulative averaging was needed to reduce statistical scatter when calculating the 'particle-to-continuum' flux.

As there is considerably less scatter associated with the measurement of macroscopic state properties, state-based coupling methods have become more popular than flux-based methods in recent years. With varying approaches to temporal coupling, these state-based methods have been applied to both steady and unsteady flows. However, strong coupling is typically used to tackle unsteady flow problems. For example, a strongly coupled state-based Euler-DSMC method was developed by Roveda *et al.* in 1998 for moving 1D shock waves [13] and then extended in 2000 to deal with 2D

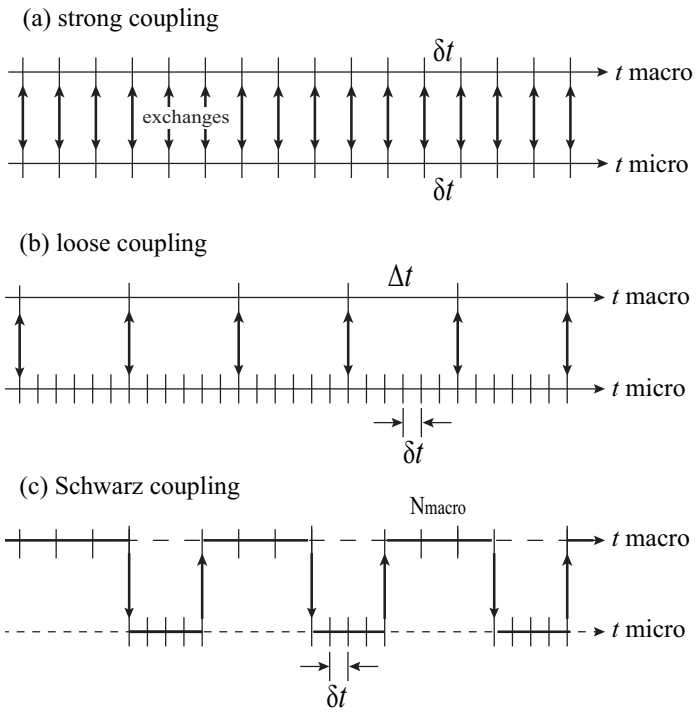

Figure 3.2: Schematic of various temporal coupling schemes used in DD hybrid methods: (a) strong coupling, (b) loose coupling, and (c) Schwarz (alternating) coupling.

unsteady slit flow [15]. Although statistical fluctuations were lessened compared with flux-based methods, the use of strong coupling meant that the averaging of the state properties in the DSMC cells was limited to the number of particles per cell. A novel approach was therefore adopted to reduce scatter in these cells: particles were effectively cloned in 'ghost cells' at no significant additional cost.

In 2003, Sun *et al.* [18] also proposed a state-based hybrid method, strongly coupling an unsteady NS solver with an information preservation (DSMC-IP) solver to simulate subsonic microscale flows. Developed by Fan and Shen [85], the DSMC-IP method is a derivation of Bird's DSMC method that is designed to reduce statistical scatter by tracking additional macroscopic information that is assigned to each particle. With reduced scatter, strong coupling could be used without instabilities. The hybrid method was demonstrated successfully for a range of microscale problems, for example, Couette flow and external flow over micro-airfoils. However, the level of non-equilibrium in these problems was fairly low. When Wang and Boyd [86] applied this scheme to hypersonic

flows (where the effects of non-equilibrium are significant), the accuracy was poor: in the simulation of 1D shock waves, the DSMC-IP solver produced an incorrect shock jump and a shock wave that was too thin. A new energy flux model helped to reduce these discrepancies somewhat, but at an increased computational cost [87].

Although a strong coupling means that the time evolution of the hybrid solution will be accurate, this type of temporal scheme is usually inefficient — advancing the continuum solver with the 'micro' time step of the particle solver is often unnecessary and wasteful. A straight-forward improvement is a 'loosely' coupled approach (referred to as the 'Continuous micro solution – Intermittent coupling' or CI scheme in [88]), where each solver adopts its own time step, and information is exchanged less frequently. This approach is shown schematically in Fig. 3.2(b). Using a much larger 'macro' time step Δt to advance the continuum solver and reducing the number of coupling exchanges can produce some computational savings, and enable us to reduce statistical scatter in the DSMC measurements with cumulative time averaging. If the macro time step is small enough, loose coupling is able to capture the evolution of the flowfield in time.

In 2009, Abbate *et al.* [21] proposed a state-based method, loosely coupling an unsteady Navier-Stokes (NS) solver to the DSMC method for unsteady flows. Both solvers were run for one 'coupling time step' that was greater than the macro NS time step and, in turn, much greater than the micro DSMC time step. Validation cases included a 1D shock tube flow and a 2D pressure-driven slit flow; the property profiles produced were shown to match those from full DSMC simulations at various time intervals, with a computational speed-up of approximately $10\times$ for the slit flow case.

More recently, in 2014 a loose state-based NS-DSMC method was developed within the OpenFOAM numerical suite by Pantazis and Rusche [24], coupling the density-based *rhoPimpleFoam* solver to the standard *dsmcFoam* solver. The creation of a hybrid solver within OpenFOAM provides opportunities for 3D unstructured meshes and parallelisation, which resulted in a highly flexible scheme that could deal with arbitrary geometries. Two unsteady problems were considered for validation: shock tube flow and flow through an orifice. The hybrid solutions were shown to match well with the equivalent full DSMC solutions at various points in time, with the computational speed-up depending on the parallel decomposition of the domain.

Garcia *et al.* [14] also used a loose temporal coupling between the continuum-fluid and DSMC solvers in the adaptive mesh and algorithm refinement (AMAR) method that was presented in 1999. AMAR is fundamentally different from traditional DD schemes (in which there is a coupling interface between two separate continuum and DSMC sub-domains) — the DSMC solver is instead introduced at the finest level of an adaptive mesh refinement (AMR) hierarchy.

The AMAR method can be considered in two parts. In the first, the Euler or NS equations are applied to the entire flowfield and the AMR procedure provides local grid refinement where needed. Starting with a coarse uniform mesh, a refinement criterion is used to determine if each cell should be divided into smaller sub-cells; this continues until the cells reach some lower size limit. In the AMR procedure, there is a coupling between the fine and coarse levels of the mesh: the quantities in the higher level grid cells are replaced by quantities that are obtained by averaging over the equivalent cells in the lower level grid. A process called 'refluxing' is then applied at the boundary of the fine grid to correct for the difference between the coarse and fine grid fluxes.

The second part of the AMAR method is an algorithm refinement. The DSMC solver now overlays the lowest level continuum cells of the AMR scheme, providing a local correction to the continuum solver. The coupling between the two solvers is similar to AMR coupling and can be considered as a combination of both state- and flux-based coupling. Each DSMC region is surrounded by 'buffer' or 'reservoir' cells (see Fig. 3.3), where a Maxwellian or CE velocity distribution is used to impose the properties of the local continuum solution. It is worth noting that, with all particles in the buffer cells being deleted at the end of each DSMC time step, this distribution must be generated at the beginning of each DSMC time step — this is a computationally costly procedure. After a number of DSMC time steps equivalent to the continuum time step, the field properties in the continuum cells are replaced by values obtained by averaging over the particles in the collocated DSMC region. The fluxes crossing the DSMC interface (averaged over the DSMC time steps to reduce scatter) are then used to correct the continuum cells adjacent to the DSMC region. To date, the AMAR has been demonstrated for 1D and 2D gas flows including the Rayleigh problem, flow past a sphere, and a moving shock problem [14, 89]. Note that the differences between this

method and the HMM-FWC method will be discussed in section 3.2.3.

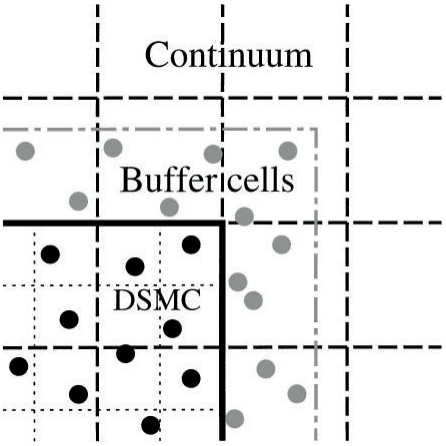

Figure 3.3: Schematic taken from [14] showing a DSMC region and its surrounding buffer cells embedded within a continuum mesh. Continuum cells (dashed lines); DSMC collision cells (dotted lines); DSMC/continuum interface (solid line); buffer cell sheath (dot-dashed line).

When coupling a costly DSMC solver with a much less expensive continuum-fluid solver, the reduction in macro time steps that is provided by loose coupling (compared with strong coupling) does not produce significant savings. The expense of continuum-particle methods comes mainly from the number of micro time steps performed to advance the particle solver, and loose coupling offers no reduction in this number. This means that strong and loose temporal coupling schemes are applicable only for short time scales; simulating the macroscopic time scales that are of interest in real engineering problems (for unsteady flows or the evolution to steady-state) is infeasible. Fortunately, for steady flows, it is possible to iterate towards steady-state instead.

A number of state-based DD hybrid methods have used the iterative Schwarz technique, which was originally adopted by Hadjiconstantinou and Patera [90] for a continuum-MD coupling. Starting from an initial guess, each Schwarz iteration lets each sub-domain reach its 'local' steady-state before transferring the state quantities in the overlap. Alternating between these local steady-state solutions, successive Schwarz iterations (typically $\mathcal{O}(20)$ [76]) are continued until the global solution converges to steady-state, i.e. the solutions from both solvers match in the overlap region. Note that in Schwarz-based continuum-DSMC methods, the continuum-fluid solver typically adopts a steady flow description and an implicit solution method; N_{macro} is the number of implicit steps required for the continuum sub-domain to reach its local converged

condition. Schwarz coupling, shown schematically in Fig. 3.2(c), is able to provide computational savings as the particle sub-domain does not require explicit integration to the global steady-state — it is generally able to reach its local steady-state in fewer micro time steps. Aktas and Aluru [16] were the first to apply a Schwarz temporal coupling scheme for rarefied gas flows. In 2002, they modelled the steady flow in microfilters using a state-based Stokes-DSMC scheme, reporting computational speed-ups ranging from 2× to 100× depending on the microfilter system.

In 2006, Wu *et al.* [19] presented a Schwarz-based NS-DSMC method that was able to take advantage of parallel computing. The method was applied to two steady-state hypersonic cases: flow over a 2D wedge, and nitrogen gas flow expanding from a 3D nozzle into a near-vacuum environment. With an unstructured mesh, the interface between the solvers was adjusted continuously. Although validation showed that the results from the wedge flow case matched well to the equivalent full DSMC solution, the hybrid simulation was found to require more computational time. Later work by Lian *et al.* [91] improved the efficiency of this method by modifying the criterion used to determine the location of the interface, enabling the method to achieve a computational speed-up of approximately 2× compared with the full DSMC simulation.

Similar to Wu's method, Schwartzentruber and Boyd [20, 92] proposed the Modular Particle-Continuum (MPC) method in 2007. Using a Schwarz approach, existing and practically unmodified DSMC and NS source codes are merged into a single modular code. While Wu *et al.* continuously exchange information between the solvers (even if the interface is not yet properly positioned), the MPC method does not update the NS boundary conditions until the interface is in the correct location. Various steady-state hypersonic flows have been simulated, including planar shock waves [20], 2D blunt body flows [93, 94], axi-symmetric blunt body flows [95], and viscous interaction flows [96]. Deschenes *et al.* have also implemented improvements in recent years, including parallelization of the code [97] and capabilities for simulating rotational and vibrational non-equilibrium [98, 99]. The computational speed-up offered by the method (compared with full DSMC simulations) has ranged from approximately 1.5× for strongly non-equilibrium flows, to around 30× for near-equilibrium flows.

Since 2007, a number of other studies have adopted similar state-based Schwarz-

type coupling for simulating steady flows in a range of applications. For example: in 2009, John and Damodaran [22] modelled the rarefied flow that occurs in the head-disk interface gap in a modern hard disk drive using a NS-DSMC method; also in 2009, Xu and Wu [23] used a zonal Euler-DSMC method to simulate unsteady high-temperature flow over a laser propelled 'lightcraft' at high altitude; in 2014, Tang *et al.* [25] investigated the nozzle and plume flowfield of small space thrusters, using axisymmetric NS-DSMC method to simulate steady nitrogen flow through a conical de-Laval nozzle as a representative test case. These studies have generally found that this type of coupling is able to match the steady-state solutions of equivalent full DSMC simulations, while obtaining some computational speed-up.

In summary, most DD continuum-DSMC methods transfer state-based information between the solvers in order to avoid the statistical scatter associated with flux-based exchanges. This state-based coupling can be strong or loose for unsteady flows, while steady flows typically use the iterative Schwarz technique. Note that both state- and flux-based continuum-MD methods have also been developed for the simulation of dense fluids [90, 100–102]. All of the DD continuum-DSMC methods that have been discussed here have imposed boundary conditions on the DSMC sub-domains through a velocity distribution that is either of Maxwellian form [12, 13, 15, 16, 19, 23, 24], or Chapman-Enskog (CE) form [14, 17, 18, 20–22, 25]; the generation of a Maxwellian distribution is common in standard DSMC codes, and the sampling of particle velocities from a CE distribution can be achieved using the algorithm of Garcia and Alder [103]. A study by Hash and Hassan [104] concluded that a Maxwellian distribution is adequate when the continuum sub-domain can be represented by the Euler equations, while a CE distribution will produce a more accurate solution when the NS equations are required.

A major challenge in the implementation of DD continuum-particle methods lies in identifying the exact regions of space and time that require particle simulation. At present, there is no single criterion used to detect the onset of non-equilibrium or continuum breakdown in dilute gases. A range of different 'breakdown parameters' have been adopted in the DD methods discussed above, including the global Knudsen number, the gradient-length local Knudsen number [105], Tsien's parameter [106], Bird's parameter [62], Tiwari's criterion [107], the parameter B [103], the criterion proposed

by Lockerby *et al.* [108], and others [109, 110]. Although the selection of a breakdown parameter is not a focus of this thesis, it may be a consideration for future work.

Generally speaking, DD hybrid methods are suitable and efficient when the continuum model is accurate across the majority of the flowfield and the molecular/particle description is required only in local regions (around singularities, or close to surfaces); these types of fluid flow problems have been labelled as 'Type A' [26]. Despite their popularity, DD methods are not useful for simulating 'Type B' problems, which require the micro solver to complement the continuum solver everywhere. This is the case when the traditional linear constitutive relations fail in the bulk of a system, for example due to non-equilibrium effects in the wake of a re-entry vehicle, or if the transport properties are unknown in an unusual gas mixture. Although less common than the DD framework, the Heterogeneous Multiscale Method (HMM) is able to deal with Type A and Type B problems, or even a mixture of the two.

3.2.2 The Heterogeneous Multiscale Method (HMM)

The Heterogeneous Multiscale Method (HMM) was first proposed in 2003 by E and Engquist [26]. The framework of the method is such that micro-resolution can be provided both close to surfaces and in the bulk of the flowfield. A continuum-fluid description (i.e. the governing conservation equations) is applied across the entire system, while the molecular/particle solver is applied in spatially distributed micro elements. As illustrated in Fig. 3.4, a 'point-wise coupling' approach is traditionally adopted: the micro elements are located at the bulk and boundary-wall nodes of the macro grid, supplying missing fluid-constitutive and/or boundary information directly to these points. At the same time as supplying this data, the micro elements must be constrained using local continuum properties obtained at the nodal points, for example the local strain-rate.

Point-wise coupling means that the HMM is effective when the spatial scales are highly separated, i.e. the variation of the macroscopic properties over the physical extent of a micro element is small. In such conditions, the HMM can provide major computational savings over a full particle simulation. However, this restriction to highly scale-separated flow systems can also be considered a disadvantage: the continuum grid spacing must be greater than the extent of the micro elements in order to avoid overlaps,

placing a limit on the minimum scale of the macroscopic variation that can be simulated [27]. The traditional HMM is therefore inefficient for fluid flow problems that possess mixed degrees of spatial scale separation.

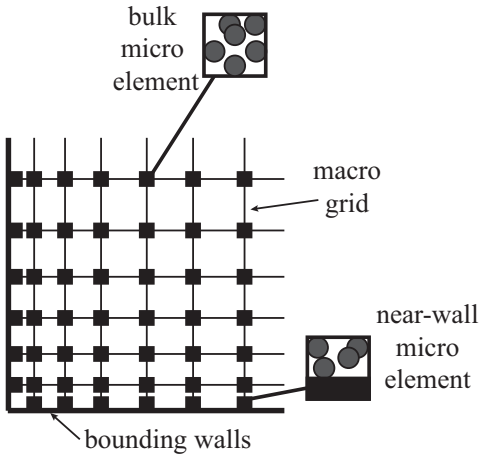

Figure 3.4: Schematic of the HMM framework with point wise coupling. Particle simulations are performed in micro elements at each node of the macro grid used by the continuum solver.

A second disadvantage of point-wise coupling is that the micro elements must be collocated with the nodes of the macro mesh. This means that, if the number of macro nodes is to be increased or reduced, then the number of micro simulations must also be increased or reduced accordingly. Linking the micro and macro resolution in this way can be restrictive as the micro and macro variations may not necessarily require the same resolution. The micro elements of the HMM typically provide information relating to the stress, which varies more slowly than macroscopic properties like velocity and density — as an example, consider the case of classical Poiseuille flow, in which the streamwise velocity profile is parabolic while the stress profile is linear; although these variations will not hold exactly at the microscale, it can be assumed that the same tendency will exist [27]. Generally, the resolution required for resolving fields like the velocity is higher than the resolution required to provide information on quantities like the stress. These disadvantages of point-wise coupling have been addressed by the 'field-wise coupling' approach of Borg *et al.* [27], which will be discussed in section 3.2.3.

The three temporal coupling schemes used by DD methods (strong, loose, and Schwarz coupling) can also be implemented for HMM hybrid methods. However, a temporal scheme that is often associated with the HMM is depicted in Fig. 3.5. Referred

to as 'Heterogeneous micro solution – Intermittent coupling' or HI in [88], this scheme is very effective for exploiting high degrees of time scale separation, i.e. for unsteady flows where the macroscopic variation in time is slow. The continuum solver (typically an unsteady description with explicit integration) takes one large macro time step Δt while the particle solver runs for a number of smaller micro time steps δt until the micro elements relax to a quasi steady-state. For highly time scale separated problems, this relaxation period t_{micro} is typically much smaller than the required macro time step, and so wasteful segments of the micro time scale can be skipped before information is exchanged. This means that, compared with a full particle simulation, significant computational savings can be obtained. Unfortunately, HI is not accurate for problems that have low or mixed time scale separation, i.e. where the macro solution varies more quickly [88].

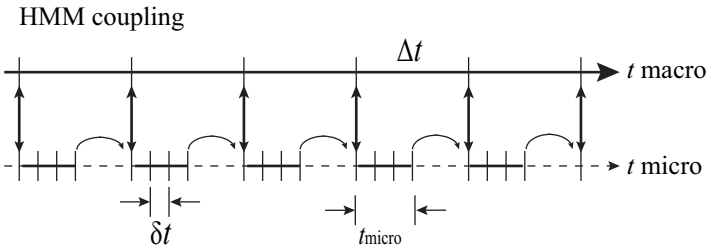

Figure 3.5: Schematic of the temporal coupling scheme (HI) often used in HMM hybrid methods.

The seamless HMM (SHMM) strategy was proposed in 2009 by E *et al.* [111] in order to overcome a potential disadvantage of the HMM algorithm: it was suggested that the HMM might be subject to a significant computational overhead if each micro element required re-initialisation after each coupling exchange. Although re-initialisation is arguably not required at all, the temporal scheme of the SHMM (known as the 'Continuous micro solution – Asynchronous coupling' or CA scheme in [88]) is an attractive alternative to HI for low and mixed levels of time scale separation. Despite exchanging data after every time step, the SHMM scheme is able to exploit time scale separation by using a different time step for each solver; each solver has its own 'clock'. The error from this asynchronous coupling is acceptable for scale-separated conditions as the micro elements remain quasi-equilibrated to the macroscopic evolution [88].

As well as assessing existing temporal schemes, Lockerby *et al.* [88] also presented

a new, highly efficient and time-accurate scheme that is essentially a generalisation of existing approaches. Known as the 'Continuous micro solution – Asynchronous Intermittent coupling' or CAI, this scheme combines the best features of loose (CI), HMM (HI), and SHMM (CA) coupling. It enables the time stepping to adapt depending on the flow behaviour: the number of micro time steps performed between exchanges is continuously adjusted. This temporal scheme could be combined with any hybrid framework to simulate unsteady flows with varying degrees of time scale separation.

To date, the traditional point-wise coupling HMM framework has been applied only for liquid flows with MD as the molecular solver. The methods developed in the literature, most notably by Ren and E [112], Yasuda and Yamamoto [113], and Asproulis *et al.* [114], focus mainly on the transfer of momentum: the micro elements are constrained using the local continuum strain-rate, and the resulting stress and/or velocity slip at the boundary-walls is passed back to the continuum nodes. These schemes differ mainly in the way that the strain-rate constraint is applied to the MD micro elements. In 2014, Yasuda and Yamamoto [115] extended their multiscale approach to consider the transfer of momentum and heat in the lubrication of a polymeric liquid with viscous heating.

In general, there has been very little progress in the development of heterogeneous hybrid methods for dilute gases, i.e. those that couple a continuum description to the DSMC method. In 2010, Kessler *et al.* [29] proposed a HMM-type framework dubbed the Coupled Multiscale Multiphysics Method (CM^3), in which an unsteady continuum description was closed using stress and heat flux fields that were obtained from a DSMC solver. Both the continuum description and the DSMC solver were, however, implemented across the entire flowfield, precluding any computational savings from spatial scale separation. With a temporal scheme similar to that of loose coupling, the CM^3 could be applied for unsteady flows (and would be accurate for mixed time scale separation). However, with no exploitation of scale separation, computational savings were obtained only by performing a smaller number of averaging ensembles in the DSMC stage of the CM^3 than were performed in the equivalent full DSMC simulation. Savings were obtained when simulating the transient start-up behaviour of a 1D Rayleigh flow — the speed-up varied from $1\times$ to $10\times$ depending on Kn.

The results from the method (both at several time instants, and the final steady-state solution) were shown to match well to the full DSMC simulation of the problem. High accuracy was also obtained for steady-state Fourier flow; unfortunately the method was found to be more expensive than the equivalent full DSMC simulation for this case. The development of a heterogeneous continuum-DSMC method that can exploit scale separation (specifically spatial scale separation) is therefore the aim of this thesis.

3.2.3 The HMM with field-wise coupling (HMM-FWC)

To overcome the disadvantages of point-wise coupling, Borg *et al.* [27] presented a modification to the traditional HMM framework in 2013, focussing on a continuum-MD coupling. Termed the 'HMM with field-wise coupling' or HMM-FWC, the framework places a continuum-fluid description across the entire domain, but each micro element now represents fields that correlate directly with a continuum sub-region, the spatial dimensions of which are identical to those of the micro element itself. The micro resolution is no longer restricted by the nodes of the macro grid: the position and size of the micro elements can be optimised independently of this grid to suit the given problem. A schematic of this HMM-FWC framework applied to a simple 1D Poiseuille flow problem is shown in Fig. 3.6(d); a comparison is made with the DD and HMM frameworks applied to the same problem. A key advantage of the HMM-FWC is its ability to simulate flow problems with mixed degrees of spatial scale separation, i.e. macroscopic properties may vary substantially across the extent of a micro element.

As well as modifying the basic framework of the HMM, the HMM-FWC also takes a somewhat different coupling approach. As discussed previously, the traditional HMM closes the governing conservation equations by applying the stresses (and sometimes heat fluxes) extracted from a micro element directly to each grid node. In the HMM-FWC, the coupling is performed via fields rather than points. The method presented in [27] assumes isothermal conditions, however, the field extracted from the micro elements is not the explicit stress — it is a constitutive stress *correction*, i.e. a correction to the linear NSF constitutive relation. These fields are then interpolated between the micro elements to provide a global correction field across the entire flowfield (unlike DD methods which provide a local correction only). Along with updated velocity bound-

ary conditions obtained from the near-wall micro elements, this global stress correction is then supplied to a 'modified' momentum conservation equation, i.e. the governing momentum conservation equation closed by the corrected constitutive relation. Essentially, the HMM-FWC takes advantage of the fact that quantities like the stress (and hence their corrections) often vary slowly in space, meaning that computational savings can be obtained by interpolating these fields between micro elements. Note that the benefit of coupling via a stress correction rather than an explicit stress field is that a correction field provides an indication of the system's departure from equilibrium.

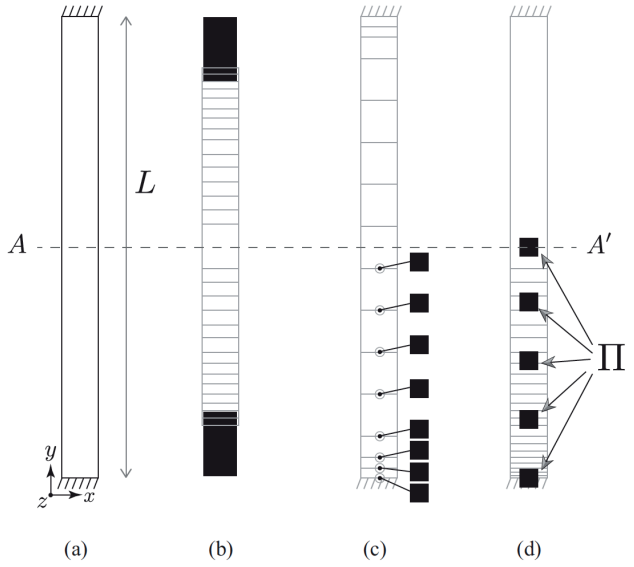

Figure 3.6: Schematic taken from [27] showing (a) 1D Poiseuille flow problem (cross-channel), and (b) the DD; (c) HMM with point-wise coupling; and (d) HMM with field-wise coupling (HMM-FWC) frameworks applied to this problem. The grey mesh indicates the macro domain, while black filled boxes indicate micro sub-domains. The A-A' is the line of symmetry.

The constraint of the micro elements is also revised. The techniques proposed in the HMM literature assume that there is no macroscopic variation across the extent of a micro element. The HMM-FWC, however, requires the imposition of a field in order to provide a direct physical correlation between the micro element and the macro domain. The local continuum solution is essentially projected onto all micro elements, i.e. the projected fields act as target conditions for each particle simulation to reproduce. The constraint procedure proposed in [27] was developed for isothermal flows with MD as the micro solver — the macroscopic velocity field (obtained from the modified momentum

conservation equation) is therefore projected onto each MD element. As an example, the constraint of a bulk micro element for a 1D Poiseuille flow problem is illustrated in Fig. 3.7. Each element is split into a 'core' region and a 'constrained' region. No constraints are applied in the core region; instead, the velocity fields imposed across the constrained regions should produce the projected velocity field $u(y)$ (see red dashed line in Fig. 3.7(b)) across the core region. As the boundary conditions of the MD elements are periodic, the velocity field imposed in the constrained regions u_i^* (see solid blue lines in Fig. 3.7(b)) is a modification of this projection that is able to satisfy periodicity. This modified velocity field u_i^* is generated by applying an artificial body force $f_{x,i}^*$ in the constrained regions, as shown in Fig. 3.7(c); this force distribution is calculated by substituting the target velocity field into the modified momentum equation. The data that is used to calculate the stress correction field is then extracted from the physically accurate core region only.

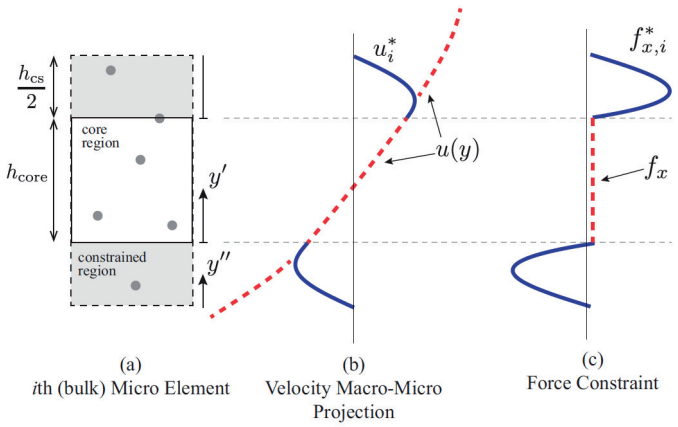

Figure 3.7: Illustration taken from [27] showing (a) a bulk micro element, (b) the projected and modified velocity fields across the periodic element, and (c) the consistent force constraint derived from the target velocity field. Dotted lines indicate unmodified projected fields, while solid lines indicate imposed fields that are modified to satisfy periodicity.

Steady, incompressible, and isothermal 1D Poiseuille flow was used to validate the HMM-FWC in [27]; an alternating Schwarz-type scheme was used to iterate towards the steady-state solutions. Two cases were considered: one using a Newtonian fluid, and the other a non-Newtonian fluid. For both cases, convergence occurred quickly (inside only 3 to 4 iterations) and the results were in close agreement with the corresponding full MD simulation results. In the Newtonian fluid case, the stress correction field constructed

from the micro elements was shown to correct for inaccuracy in the assumed viscosity model. Similarly, in the non-Newtonian fluid case, the stress correction was able to provide the unknown fluid-constitutive behaviour. A computational speed-up of 74× was achieved for the Newtonian fluid case (which required only one near-wall micro element for this axisymmetric problem) while a speed-up of 66× was obtained for the non-Newtonian case (which required one near-wall element and one bulk element).

At first glance, the HMM-FWC approach might seem similar to the AMAR method. There are, however, some important differences. First, AMAR can be considered a DD approach as it provides only a local 'correction' to the overall solution; on the other hand, the HMM-FWC provides a global correction that covers the entire flowfield. Secondly, the procedure for constraint of the particle sub-domains is different. The AMAR method (which was implemented for dilute gases) imposes the local continuum properties in buffer cells that surround each DSMC sub-domain by generating a particle distribution based on these properties; this approach is, however, expensive as all of the particles in the buffer cells are inserted and deleted at every DSMC time step. Although the HMM-FWC has been implemented with MD micro elements, the general approach to constraint is based on enabling the core region of the element to relax gradually in space and time — it can therefore be considered a more natural and less expensive approach. The AMAR coupling also includes flux-based exchanges, which are not performed in the HMM-FWC. Finally, the AMAR uses a loose coupling to deal with unsteady flows, while the HMM-FWC has been applied only to steady flows using a Schwarz-type coupling.

As the HMM-FWC can deal with various degrees of spatial scale separation, it can be considered the most general hybrid framework. For this reason, the continuum-DSMC method that will be presented in the following chapter of this thesis uses a HMM-FWC framework. The constraint of the micro elements will be modified to suit the use of the DSMC method, and the coupling between the solvers will be extended to deal with mass, momentum, and heat transport such that compressible and non-isothermal flows can be simulated. Only steady flows will be considered and so an iterative Schwarz-type coupling scheme will be sufficient; future development could involve extension to time-scale separated problems by using the HI, CA, or CAI schemes.

3.2.4 The Internal-flow Multiscale Method (IMM)

The final framework that is discussed in this chapter can be considered a specific class of the HMM-FWC. Presented by Borg *et al.* [28] in 2013 for a continuum-MD coupling, the Internal-flow Multiscale Method (IMM) is tailored to simulate flows through long micro/nanochannels. Such channels are common in emerging miniaturised technologies, and are generally of high aspect-ratio, i.e. the length in the streamwise direction is likely to be several orders of magnitude larger than the dimensions of the cross-section. Fluid flows in this type of geometry cannot be classed as Type A or Type B, and so have been termed 'Type C' problems [116]. As the majority (if not all) of the flow can be considered near-wall, a DD approach would require the micro solver to occupy a large portion of the computational domain, as indicated in Fig. 3.8(b); this would result in few savings and would severely restrict the length of channel that could be studied. The traditional point-wise coupling HMM would also be inefficient for such highly-confined channel flows as the micro elements would overlap; see Fig. 3.8(c).

The IMM can instead obtain computational savings by exploiting the uni-directional length scale separation that exists in these high-aspect-ratio flows: the molecular processes occurring transverse to the flow are only loosely coupled to the gradual variation of the macroscopic properties in the streamwise direction. Like the other heterogeneous frameworks, the IMM applies a continuum description over the entire channel. The molecular/particle solver is then implemented in micro elements that are distributed along the channel length, covering the entire cross-section as shown in Fig. 3.8(d).

The coupling exchange of the IMM differs from that of the HMM and the HMM-FWC as the boundary and constitutive information is supplied to the continuum description indirectly. The complexities of imposing strain-rates/velocity fields on the micro elements and extracting the resulting stresses are avoided. Assuming an isothermal, incompressible, and steady-state flow, the IMM coupling is based on imposing pressure gradients (or pressure drops) across the micro elements, and using the resulting mass flux to correct the continuum description. As the boundaries of the micro element are periodic, these pressure drops are imposed by simply applying a body force to the molecules. Using a Schwarz-type iterative approach, the iterations continue until the mass and momentum flux are both conserved across the system.

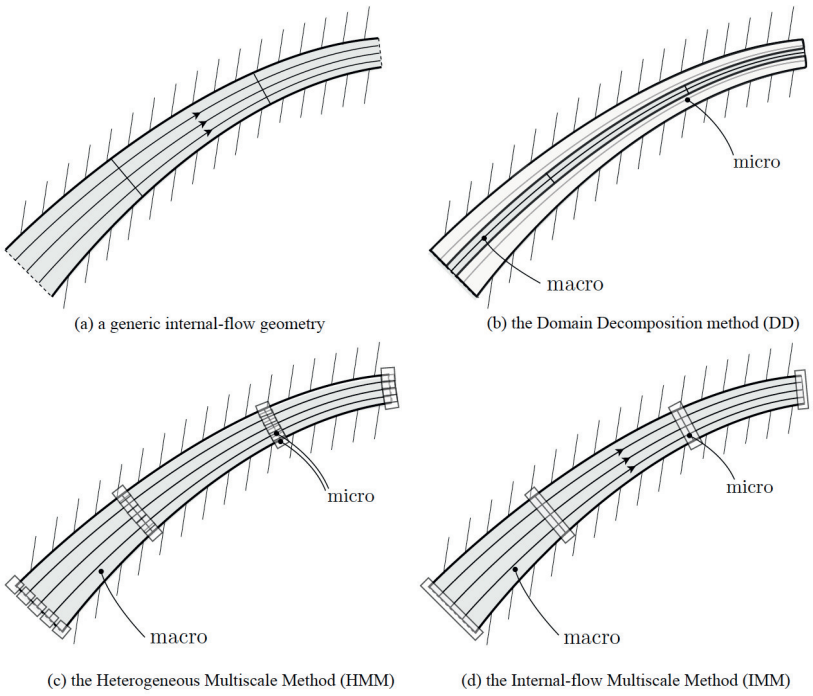

Figure 3.8: Schematic taken from [30] of different types of hybrid frameworks applied to (a) a generic internal-flow problem: (b) the DD, (c) the HMM, and (d) the IMM.

The IMM was originally demonstrated on force-driven flow through an axially-periodic converging/diverging nanochannel. Later in 2013, Borg *et al.* [117] also considered its application to network-type systems that can be decomposed into high-aspect-ratio channel components and junction components — computational savings can then be obtained by modelling the high-aspect-ratio channels via a series of smaller micro elements, while the junction components (for example a reservoir or a bend) are simulated entirely by the micro solver. Attention was restricted to serial networks (i.e. where each component has only one inlet and one outlet) in the form of pressure-driven flow through a nanochannel that connects two reservoirs. For both of these applications, the convergence of the method was found to be very quick, occurring inside 2 to 4 iterations. The mass flow rate/mass flux results showed good quantitive agreement with equivalent full MD simulations, with only small errors in the range of $1 - 5\%$ depending on the case and the number of micro elements used. There were, however, discrepancies between the streamwise pressure and density variations predicted by the IMM and those predicted by the full MD simulations. The authors attribute this error

to the effects of compressibility — even at low Mach number in isothermal conditions, compressibility can have significant effects in high-aspect-ratio geometries due the combination of high viscous forces and the long channel lengths. Computational speed-ups ranging from approximately 3× to 12× were achieved, with larger speed-ups likely for longer (and more realistic) channels.

In 2014, Stephenson *et al.* [118] extended the IMM technique to general micro/nano network geometries of arbitrary size and complexity, where components can have any number of inlets and outlets. The MD micro elements and junction components were required to satisfy the pressure variation and mass conservation, both locally and globally. A density correction was also implemented to account for compressibility effects: the fluid density in each component was adjusted according to an equation of state (based on the applied boundary pressures). The flow of liquid argon through a bifurcating channel geometry was considered for (a) a bifurcating case, with one branch splitting into two; and (b) a mixing case, with two branches joining into one. Convergence of the IMM occurred inside 3 iterations for both cases and, comparing against full MD simulations, the mean errors in the mass flow rate and pressure variations were below 4%. Speed-ups of approximately 2× were obtained.

Unlike the HMM-FWC of [27], there have been a number of extensions to the original IMM of [28, 117]. In 2013, Patronis *et al.* [30] considered the application of the IMM to rarefied dilute gases, and also extended the coupling technique to deal with compressible flows. Replacing MD with the DSMC method in the micro elements did not require significant modifications to the IMM algorithm — the largest modifications came instead from the compressibility requirements. The incompressible IMM of [28, 117] depended on calculating the pressure drop that should be applied across each micro element in order to represent the pressure difference across the full system (while also ensuring that mass conservation was satisfied). In addition to obtaining this pressure drop, the compressible IMM must also obtain the fluid density that should be used to initialise each micro element. The low-speed flow of a rarefied gas through three high-aspect-ratio geometries was considered for validation: streamwise-periodic converging-diverging microchannel flow; non-periodic pressure-driven microchannel flow around a corner; and eccentric cylindrical micro Couette flow. For all three, the compressible

IMM was able to produce flow predictions that matched extremely well to the equivalent full DSMC simulations while providing speed-ups of 6×, 50×, and 300×, respectively.

A further extension of the IMM to non-isothermal flows was presented by Patronis and Lockerby [31] in 2014. The method was applied to investigate thermal creep/ transpiration, and its use in high-aspect-ratio Knudsen-compressor channels. Thermal creep is a rarefied gas effect, where the presence of a streamwise temperature gradient results in the development of a slip flow at a surface; this flow moves, counter-intuitively, from low temperatures to higher temperatures. Solid-state Knudsen-compressors operate by exploiting this behaviour. For this type of application, the compressible IMM must predict the effects of the applied streamwise temperature variation on both the mass flow rate and the streamwise pressure and density variations. The reported modification to the IMM coupling strategy was not substantial: as well as applying the correct pressure drop and density to each micro element, the temperature and the temperature gradient dictated by the wall boundary conditions for that element were also imposed. As examples, temperature profiles were applied to straight, converging-diverging, and diverging-converging channel geometries. For all cases, the IMM was able to provide flow predictions that matched very well to the equivalent full DSMC simulations, while achieving major computational savings. Note that, despite being effective for the IMM applied to these thermal transpiration flows, this coupling strategy is not general; the extension of the HMM-FWC to non-isothermal flows will require coupling via the conservation of energy. It should also be noted that the micro solver used in [31] was not the traditional DSMC method — a low-variance deviational simulation Monte Carlo (LVDSMC) method was applied in the micro elements instead.

The use of LVDSMC comes from the desire to avoid the statistical scatter that can be problematic in DSMC, particularly for low-speed flows. The LVDSMC achieves variance reduction by splitting the velocity distribution into an equilibrium distribution that is represented analytically, and a deviational distribution that is represented by particle simulation. Using the same number of particles that would have been used to model the full velocity distribution, the variance in the deviational distribution is hugely reduced [31]. In the non-isothermal IMM algorithm, the imposition of the boundary conditions on each micro element was tailored to the capabilities of the LVDSMC

method: the streamwise gradients in pressure and temperature were both represented (independently of each other) by an effective body force on the particles. This meant that the micro elements remained periodic, avoiding the complexity and computational expense of non-periodic boundaries. Although the LVDSMC method is suitable for this particular application, it is accurate only if the departure from equilibrium is small. The LVDSMC source code [119] is also subject to geometric limitations at present.

The latest extension to the IMM is the 'unsteady-IMM' that was presented by Borg *et al.* [32] in 2015. Based on a continuum-MD coupling, this method adopts the CAI temporal scheme proposed in [88]. By exploiting varying levels of both length and time scale separation, the unsteady-IMM was able to achieve orders of magnitude computational savings (compared with full MD simulations) when applied to a converging-diverging nanochannel geometry, and the level of accuracy was excellent in most cases. The degree of time scale separation in a given case was found to determine not only the computational savings achieved, but also the sensitivity of the method's accuracy to the macro-micro time stepping.

3.3 Summary

This review has highlighted that a large majority of the continuum-particle methods in the literature utilise a DD framework, with the exchange of information between the solvers being typically state-based to avoid the statistical scatter that is inherent in flux-based exchanges. Despite their popularity, DD hybrid methods are effective only for Type A problems, where the particle description is required only in localised regions (e.g. close to bounding walls). The framework of the HMM provides an interesting alternative to DD as it is suitable for Type B problems, in which the micro solver is required to complement the continuum solver everywhere due to the failure of the traditional linear constitutive relations. However, the point-wise coupling approach of the HMM means that the micro resolution (i.e. the position and size of the micro elements) is restricted by the need for collocation with the nodes of the macro grid; it also restricts the use of the HMM to flows where the spatial scales are highly separated.

The development of the field-wise coupling (HMM-FWC) approach was the result of an effort to overcome these restrictions. In this framework, the micro resolution is

not restricted by the nodes of the macro grid — the micro elements can be any size and can be placed at any location in the flowfield. A constitutive correction field is obtained from all micro elements, and is then interpolated between adjacent elements to provide a global correction. Unlike the traditional HMM, the HMM-FWC is then able to simulate flows with mixed degrees of spatial scale separation; it can therefore be considered a more general heterogeneous approach. The IMM can be viewed as a type of HMM-FWC with a framework and coupling strategy specifically tailored to Type C flows through high-aspect-ratio channels (where large length scale separation exists in the streamwise direction). Both the HMM-FWC and the IMM were originally presented with a continuum-MD coupling for the simulation of liquids, and assumed steady, incompressible, and isothermal flow. To date, the IMM has seen more development than the HMM-FWC, with extensions to rarefied gases (using DSMC as the micro solver), and to compressible, non-isothermal, and unsteady flows.

In this thesis we will develop a continuum-DSMC method within a HMM-FWC framework, and will extend the coupling strategy to deal with compressible and non-isothermal flows. An extension to unsteady flows will not be considered but could be a subject of future work.

Chapter 4

A new heterogeneous continuum-DSMC method

In this chapter, a heterogeneous hybrid method is developed for dilute gases, coupling a continuum-fluid description to a DSMC solver. An HMM-FWC framework is adopted, and the coupling strategy is extended from that of the original HMM-FWC (which considered the transport of mass and momentum) to consider the transport of energy. A new procedure for constraining the micro elements is implemented: the imposition of the local continuum property fields is now tailored to the use of a DSMC micro solver. It is also worth noting that, while the original HMM-FWC was tested only on a 1D flow problem, this method will be validated and tested on 1D and 2D flow problems in subsequent chapters.

The extension of the coupling strategy means that, unlike a large majority of hybrid methods in the literature, our method is able to deal with non-isothermal flow problems; it is also able to model compressible flow. As was discussed in section 3.2.3, the HMM-FWC framework is adopted because it is the most general heterogeneous framework and has the potential to tackle flow problems that are beyond the reach of the other hybrid frameworks. Unlike the CM^3 of [29], it is able to exploit spatial scale separation. It is suitable for Type A problems where local non-equilibrium occurs, for example, close to a surface. It is also suitable for Type B scenarios where the traditional linear constitutive relations fail in the bulk of a system, meaning it is able to capture non-equilibrium behaviour in the bulk, or compensate for unknown transport properties (for

example, in an usual gas mixture). The key benefit of the HMM-FWC is that it is not subject to the restrictions of the HMM and the IMM: the micro elements can be any size and can be placed independently of the continuum grid, as illustrated in Fig. 4.1; flow problems with mixed degrees of spatial scale separation can then be simulated.

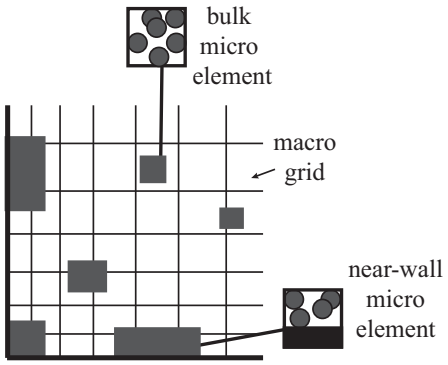

Figure 4.1: Schematic of the HMM with field wise coupling (HMM-FWC). DSMC simulations are performed in micro elements that are positioned independently of the macro grid used by the continuum solver.

An example of a flow problem that exhibits mixed degrees of spatial scale separation is pressure-driven flow through a microchannel that is connecting two large reservoirs; see Fig. 4.2. This system might be simulated to investigate inlet and outlet losses, for example. If the HMM-FWC was implemented, larger micro elements would be placed in the reservoirs where the spatial scale separation is low, while smaller elements would be distributed along the length of the micro channel where the streamwise scale separation is high. Note that this problem could also be simulated by combining the HMM-FWC with the IMM, i.e. the HMM-FWC method would be applied in the reservoirs and at the inlet and outlets of the channel, and the IMM would be applied along the channel length.

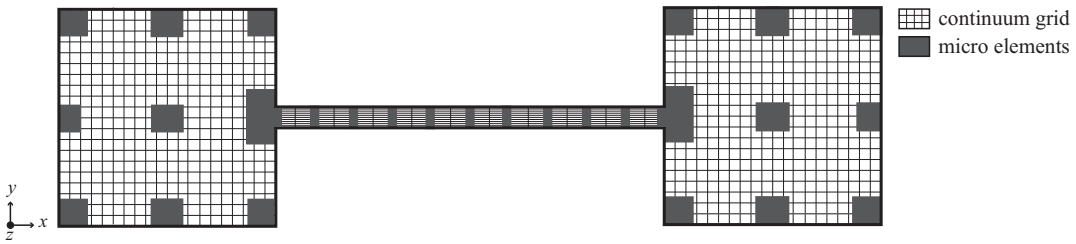

Figure 4.2: Schematic of the HMM-FWC framework applied to flow through a microchannel that is connecting two reservoirs. The continuum-fluid grid is applied across the entire system and the particle simulations are performed in the micro elements.

Another flow problem that could exhibit mixed degrees of spatial scale separation is the flow of a rarefied gas through a narrow microscale crack. As mentioned in Chapter 1, the DD, HMM, and IMM frameworks would not be able to exploit high scale separation (i.e. slow variation of the flow properties) over the large dimension of the flow cross-section. The application of our HMM-FWC continuum-DSMC method to this type of flow problem will be discussed in Chapter 7 — as the method can deal with heat transport, the presence of a temperature difference between the walls of the crack will be considered.

4.1 Hybrid methodology

The *dsmcFoamStrath* solver is used in the micro elements of this hybrid method, and for the full DSMC simulations that are required to validate the method; several features of this DSMC code (its measurement and state controller capabilities in particular) are useful for implementing the hybrid method's coupling algorithm. Note that, for all of the test cases that will be simulated in the later chapters of this thesis, the micro elements of the hybrid method are small enough that each micro DSMC simulation is performed in serial (i.e. on a single processor). All of the full DSMC simulations are, however, performed in parallel (i.e. they are split across a number of processors).

The steady-state continuum conservation equations can be simplified somewhat for the 1D and 2D validation problems that will be considered in this thesis. Finite difference approximations are adequate, and can be solved using simple matrix manipulation in MATLAB. In future, the simulation of more complex 2D and 3D cases could utilise a CFD solver from the OpenFOAM toolbox; for example, *rhoSimpleFoam* is suitable for steady-state, compressible flows.

At present, the algorithm of the hybrid method is also implemented using a MATLAB script — this handles all the requirements of the coupling, including the creation of the micro DSMC simulations, the execution of both solvers, and the exchange of information between the solvers. The migration of the entire hybrid algorithm into OpenFOAM is, however, an aim of future work.

As attention is restricted to steady-state flow problems in this thesis, an alternating Schwarz-type scheme is adopted; see Fig. 3.2(c). However, the use of simplified finite

difference approximations means that the solution of the continuum equations requires only one implicit step (i.e. $N_{\text{macro}} = 1$). It is worth noting that DD methods generally require that a system is initialised uniformly across the flowfield, with Schwarz coupling then converging in $\mathcal{O}(20)$ iterations; the HMM-FWC approach, on the other hand, allows initialisation based on some initial guess about the flow variation across the system, and so the correct solution can be reached in fewer iterations (usually <10). As mentioned in section 3.2.3, an extension to unsteady flows could be a focus of future work, and would involve the use of a HI, CA, or CAI scheme.

The general methodology of our HMM-FWC continuum-DSMC method is as follows. A 'modified' continuum-fluid description is applied across the entire flowfield; the modification comes from closure of the governing conservation equations with 'corrected' constitutive relations, i.e. the linear NSF relations augmented by correction fields. At the start of the simulation, these constitutive correction fields are assumed equal to zero, and the conventional no slip/jump boundary conditions are imposed. Micro elements are dispersed across the system with the arrangement depending on the flow problem[1]. Based on the local continuum solution, constraints are applied to the micro elements and DSMC is performed. After steady-state is reached and a subsequent cumulative averaging period is complete, the macroscopic property fields are extracted from each micro element and used to compute the local correction fields across that element; the global correction fields (across the entire flowfield) are then approximated by interpolating between adjacent micro elements. The macroscopic properties of the gas at the bounding walls are also extracted from the near-wall micro elements and interpolated to provide boundary conditions at all bounding walls. Using these global correction fields and updated boundary conditions to solve the modified continuum equations then provides updated or 'corrected' continuum property fields — these are used to constrain the micro elements once more, and this process is repeated. With continuing iterations, the corrected continuum solution should converge towards a full DSMC solution of the problem. The two-way coupling exchange performed in each iteration is summarised in the schematic of Fig. 4.3.

[1] In this thesis, the placement of the micro elements is based on the author's judgement (and the knowledge that non-equilibrium flow usually occurs close to bounding walls). Future work could include the selection or definition of a breakdown parameter for identification of localised spatial regions that require micro resolution.

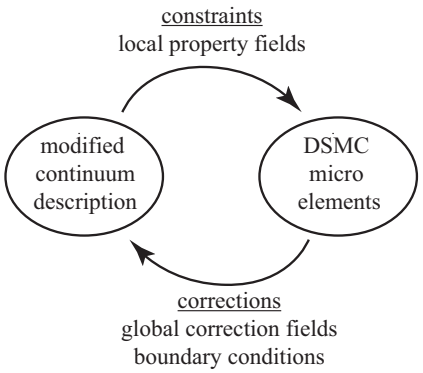

Figure 4.3: Schematic of the coupling exchange performed by the hybrid method in each iteration.

4.2 Coupling strategy

The coupling strategy is now discussed in detail; note that this discussion applies to 1D, 2D, and 3D problems. As mentioned previously, the transport of energy is now considered, and the constraint of the micro elements is now tailored to the use of a DSMC solver.

4.2.1 Macro-to-micro coupling: constraining the micro elements

The corrections to the constitutive relations must be extracted from micro elements that properly represent the local conditions in the macro domain. For dilute gases, this is achieved when the correct particle distribution is imposed at the boundaries of the micro element. However, the particle distribution required at the boundaries cannot be extracted directly from the macroscopic continuum solution — such detail is not available from a continuum-fluid description. This problem is circumvented here by introducing an artificial 'relaxation' region around each micro element that facilitates a natural relaxation to the correct particle distribution (and macroscopic state) within the core of the element. It is not essential that the gas state in the relaxation region accurately represents the conditions in the corresponding macro domain; its sole purpose is to develop the boundary conditions needed for the core region of the element. Sampling of the property fields is then performed only in this core 'sampling region'. A schematic of a 3D bulk micro element is shown in Fig. 4.4. Note that these sampling and relaxation regions can be defined easily in the *dsmcFoamStrath* code by generating zones on the DSMC computational mesh.

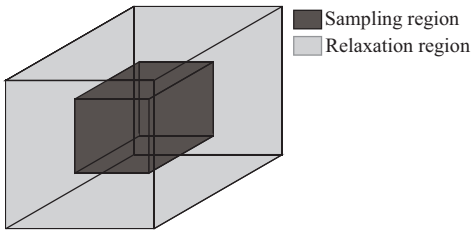

Figure 4.4: Schematic of a 3D bulk micro element, showing the sampling and relaxation regions.

The imposition of boundary conditions is now tailored to dilute gases (with DSMC as the micro solver) and is achieved in two stages. First, the local continuum velocity, temperature, and density *fields* (i.e. the spatial variation of these properties) are imposed throughout the relaxation region. Differing from the body forcing approach of the original HMM-FWC, this is done by implementing state controllers (as described in section 2.2.2) for the duration of the micro DSMC simulation. Secondly, although a particle distribution is not applied directly at the boundaries of the sampling region, a particle distribution is applied at the outer boundaries of the relaxation region. This applied distribution replaces the periodic boundaries implemented in the original HMM-FWC, removing the need for the imposed property fields to satisfy periodicity. In this thesis, local Maxwellian distributions are applied for simplicity[2]. The relaxation region must then be large enough that the particle state will relax fully (via particle collisions) across its extent — the particle distribution in the sampling region will then be dictated solely by the applied continuum state, with no effects from the imposed equilibrium distribution. When the hybrid solution converges, the artificiality of the relaxation region will dissolve seamlessly into the true particle distribution and fluid state in the core sampling region.

Force-driven flows are modelled in pure DSMC simulations by applying an external forcing (i.e. acceleration) to all particles in the system; although the acceleration is applied continuously, particle-wall interactions mean that the macroscopic velocity will reach a constant steady-state. In an equivalent hybrid simulation, we apply the same acceleration to all of the particles in the micro elements (in both the relaxation and sampling regions), in addition to the micro constraints discussed above.

[2]Local Chapman-Enskog (CE) distributions could also be applied using the approach of Stephani *et al.* [120] — as this distribution incorporates a perturbation from equilibrium, the extent of the required relaxation region may be reduced somewhat.

In summary, the sampling region of a micro element is surrounded by a relaxation region: the local continuum property fields are imposed in this relaxation region, and a particle distribution is applied at its outer boundaries; if required, an external acceleration is also applied to all particles in the element. The sampling region of a near-wall micro element must, however, be adjacent to the wall itself in order to capture the non-equilibrium behaviour that exists at bounding walls. Figure 4.5 shows bulk and near-wall micro elements in an example computational set-up for a 2D problem.

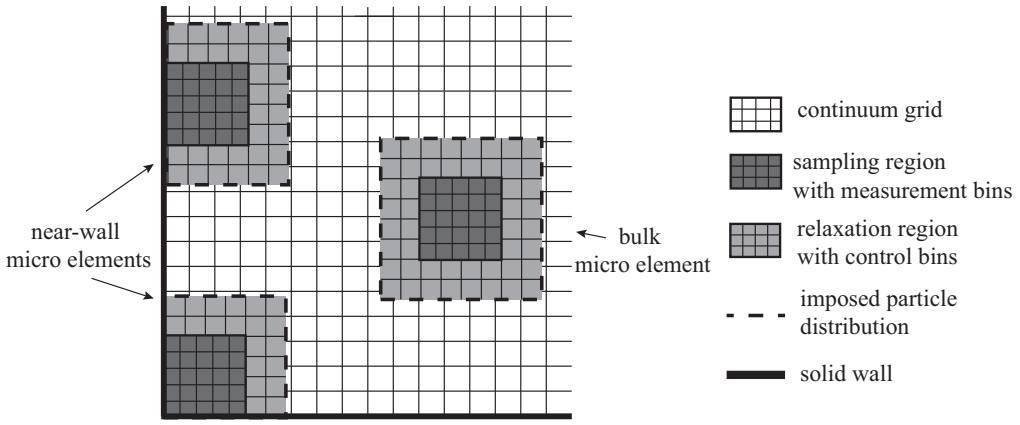

Figure 4.5: Schematic of an example computational domain for a 2D problem showing bulk and near-wall micro elements. The control and measurement bins are independent of the DSMC computational cells, and can also be independent of the continuum mesh.

Our coupling strategy requires that the *spatial variation* of the flow properties is imposed on and extracted from regions of the DSMC micro elements. As was discussed in section 2.2.2, macroscopic fields can be imposed on and extracted from a flowfield in *dsmcFoamStrath* by using the MOB technique. Each relaxation region is therefore divided into a grid of control bins; state controllers are used to impose the continuum properties in each control bin, creating the desired spatial variation across the region. Similarly, each sampling region is divided into a grid of measurement bins and the property fields are extracted by averaging over the particles in each bin according to Eqs. (2.15) – (2.24). In the near-wall micro elements, boundary information is also extracted from the wall-adjacent faces of the measurement bins by summing over all particles that strike the surface according to Eqs. (2.25) – (2.28). The key advantage of using the MOB here is that the measurement and control bins are completely independent of the computational mesh used by the DSMC algorithm. As mentioned in

section 2.2.2, this enables adjustment of the macroscopic resolution of the fields that are extracted/imposed, and provides the ability to control the level of statistical scatter in the measurements without affecting the accuracy of the DSMC computations. The measurement and control bins can also be independent of each other, and of the continuum grid — if the measurement and control bins are not collocated with the cells of the continuum mesh, then data can simply be interpolated between the meshes using MATLAB interpolation functions.

Sufficient micro resolution is crucial to obtaining an accurate result from the hybrid method, and the adequate micro element configuration (i.e. the number, location, and size of the micro elements) will vary from problem to problem. However, for a given number of micro elements at predefined locations, we enable the spatial extents of both the sampling and relaxation regions to adapt dynamically depending on the local flow by defining these extents as some number of local mean free paths λ_l — the appropriate number for each region will be studied in Chapter 5. At each iteration, the dimensions of each element are then set depending on λ_l: in the first iteration, λ_l is assumed equal to the global mean free path, λ_{gl}; in subsequent iterations, λ_l is updated for each micro element using Eqs. (2.8) and (2.9), where T and n are the are spatial averages measured in the sampling region of the element in the previous iteration.

With the element extents potentially changing at each iteration, all of the micro DSMC simulations are initialised at equilibrium (by sampling the particle velocities from a Maxwellian distribution) at the start of each iteration. For consistency between the micro and macro domains, the velocity, temperature, and density for initialisation are obtained from the local continuum solution: continuum values of these properties are averaged over a sub-region of the macro domain that corresponds to the element.

4.2.2 Micro-to-macro coupling: correcting the continuum description

The ease in obtaining macroscopic properties from particle information means that micro-to-macro coupling is less challenging than macro-to-micro coupling. However, the data transferred from the micro elements to the continuum description must not contain too much statistical scatter, as this could lead to instability of the coupling algorithm. Each micro DSMC simulation (and full DSMC validation simulation) is therefore

performed in two stages: a transient period enables the flow to reach steady-state, and a longer cumulative averaging period reduces scatter in the measured properties[3].

The steady-state conservation equations of fluid mechanics form the basis of our hybrid method, i.e.

$$\nabla \cdot \rho \mathbf{u} = 0, \tag{4.1}$$

$$\rho \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla \cdot p\mathbf{I} + \nabla \cdot \boldsymbol{\tau} + \mathbf{f}, \tag{4.2}$$

and,

$$\nabla \cdot (\rho e \mathbf{u}) = -\nabla \cdot (p\mathbf{I} \cdot \mathbf{u}) + \nabla \cdot (\boldsymbol{\tau} \cdot \mathbf{u}) - \nabla \cdot \mathbf{q}, \tag{4.3}$$

where $\mathbf{u}$ is the velocity vector, $\boldsymbol{\tau}$ is the viscous stress tensor, $\mathbf{f}$ is an external force per unit volume, e is the specific energy, and $\mathbf{q}$ is the heat flux vector. For most conventional flow problems, the linear NSF constitutive relations are used to close these continuum equations. However, the interest here lies in flows where these linear relations fail. In our HMM-FWC hybrid method, the constitutive relations used for closure are the linear NSF relations augmented by 'correction fields', i.e.

$$\boldsymbol{\tau} = \mu(\nabla \mathbf{u}) + \mu(\nabla \mathbf{u})^T - \frac{2}{3}\mu(\nabla \cdot \mathbf{u})\mathbf{I} + \boldsymbol{\Omega}, \tag{4.4}$$

and,

$$\mathbf{q} = -\kappa \nabla T + \boldsymbol{\Phi}, \tag{4.5}$$

where $\boldsymbol{\Omega}$ is a stress correction tensor and $\boldsymbol{\Phi}$ is a heat flux correction vector, μ is the gas dynamic viscosity, and κ is the gas thermal conductivity. Closing Eqs. (4.2) and (4.3) with these 'corrected' constitutive relations gives,

$$\begin{aligned} \rho \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla \cdot p\mathbf{I} + \nabla \cdot \mu(\nabla \mathbf{u}) + \nabla \cdot \mu(\nabla \mathbf{u})^T \\ -\frac{2}{3}\nabla \cdot [\mu(\nabla \cdot \mathbf{u})\mathbf{I}] + \nabla \cdot \boldsymbol{\Omega} + \mathbf{f}, \end{aligned} \tag{4.6}$$

[3]Ensemble averaging could be used to reduce scatter if transient flows are considered in future.

and,

$$\begin{aligned}\nabla \cdot (\rho e \mathbf{u}) = -\nabla \cdot (p\mathbf{I} \cdot \mathbf{u}) + \nabla \cdot [\mu(\nabla \mathbf{u}) \cdot \mathbf{u}] + \nabla \cdot [\mu(\nabla \mathbf{u})^T \cdot \mathbf{u}] \\ -\frac{2}{3}\nabla \cdot [\mu(\nabla \cdot \mathbf{u})\mathbf{I} \cdot \mathbf{u}] + \nabla \cdot [\mathbf{\Omega} \cdot \mathbf{u}] + \nabla \cdot (\kappa \nabla T) - \nabla \cdot \mathbf{\Phi},\end{aligned} \tag{4.7}$$

respectively. Along with the gas equation of state,

$$p = \rho R T, \tag{4.8}$$

Eqs. (4.1), (4.6), and (4.7) can then be used to describe the entire flowfield. As a starting point, the correction fields $\mathbf{\Omega}$ and $\mathbf{\Phi}$ are initially assumed equal to zero; in subsequent iterations, they are supplied by the distributed DSMC micro elements, incorporating the effects of non-equilibrium and/or any inaccuracies in the assumed transport properties, μ and κ.

The extraction of $\mathbf{\Omega}$ and $\mathbf{\Phi}$ from the micro elements (both near-wall and bulk) is performed in two steps. In the first step, the *local* correction fields across each sampling region are computed: the corrections in each measurement bin and at each wall boundary are calculated from the time-averaged flow properties that were measured in that bin or at that wall boundary. The stress correction is calculated according to Eq. (4.4) based on the measured stress and the velocity gradient; similarly, the heat flux correction is calculated from Eq. (4.5) using the measured heat flux and the temperature gradient. Note that the property gradients in each bin (and at each wall boundary) are approximated using finite difference representations that are based on the properties in adjacent bins (or the adjacent wall boundary); for example, the temperature gradient in a bin b is approximated as $\frac{dT}{dx}_b = \frac{(T_{b+1} - T_{b-1})}{2\delta x}$, where δx is the bin width in the x-direction. The second step is then to obtain approximations of the *global* correction fields across the entire domain, for example by interpolating between adjacent local correction fields. In this thesis, linear interpolation is initially adopted due to its simplicity. The use of higher order interpolations and global fitting approaches will be considered in Chapter 6.

In addition to the constitutive correction fields, the micro elements also provide the continuum description with updated boundary information. The properties (i.e.

the velocity, temperature, and pressure) of the gas at the wall are extracted from the wall-adjacent faces of the measurement bins in each near-wall sampling region — this local boundary information is then interpolated (or extrapolated) to produce estimates at all bounding walls. To maintain simplicity, this interpolation/extrapolation is typically linear. Therefore, when simulating 2D and 3D problems, the configuration of the near-wall micro elements will be dictated by the need to construct a reasonable approximation of the boundary properties with this type of interpolation.

4.3 The iterative coupling algorithm

As mentioned above, the hybrid method is iterative for steady flows. The method's algorithm is therefore executed according to the following steps:

(0) Assuming no constitutive corrections (i.e. $\mathbf{\Omega} = \mathbf{\Phi} = 0$) and no slip/jump at the bounding walls, solve the continuum equations (4.1), (4.6), and (4.7), and the gas equation of state (4.8), to obtain an initial NSF velocity field $\mathbf{u}_{\text{NSF}}$, temperature field T_{NSF}, pressure field p_{NSF}, and density field ρ_{NSF} across the entire system.

(1) Initialise and constrain each DSMC micro element:

   (a) Impose the local continuum property fields (i.e. the velocity, temperature, and density) across the relaxation region by implementing state controllers in each control bin.

   (b) Apply a Maxwellian particle distribution at the outer boundaries of the relaxation region based on the local continuum temperature and velocity.

   (c) If a condition of the flow problem, apply an external acceleration to all particles in the element.

(2) Execute the DSMC method in each micro element. When steady-state is reached, perform cumulative averaging of the macroscopic measurements for some averaging period.

(3) When all DSMC micro simulations are complete, extract the time-averaged flow properties from the measurement bins of each sampling region. Additionally, extract the gas properties at the bounding walls from the wall-adjacent bin faces.

(4) Calculate the velocity and temperature gradients in each measurement bin and at each bounding wall using (here) finite difference approximations based on the properties in the adjacent bins or at the adjacent bounding wall.

Using the velocity gradient and stress from all measurement bins/bounding walls, compute the local stress correction field across each sampling region via Eq. (4.4). Similarly, using the temperature gradient and heat flux from all measurement bins/bounding walls, compute the local heat flux correction field across each sampling region via Eq. (4.5).

(5) Based on the local correction fields, perform interpolations (or global fitting) to approximate the global correction fields $\mathbf{\Omega}$ and $\mathbf{\Phi}$ across the entire flowfield. Interpolate (or extrapolate) linearly between the local gas boundary properties obtained in Step (3) in order to estimate the gas properties at all bounding walls.

(6) With these global correction fields and this updated boundary information, solve the continuum equations (4.1), (4.6), and (4.7), and the gas equation of state (4.8), to obtain a new, corrected velocity field $\mathbf{u}_{\text{corr}}$, temperature field T_{corr}, pressure field p_{corr}, and density field ρ_{corr} across the system.

(7) Using this corrected solution to constrain the micro elements, repeat from Step (1) until the corrected property fields do not change between iterations to within user-defined tolerances. The final converged solution should then agree well with a full DSMC solution of the same problem.

4.4 Summary

A heterogeneous hybrid method has been proposed, in which the DSMC method complements a continuum-fluid description, i.e. provides missing boundary and constitutive information. Using field-wise coupling, the key advantage of this new method over existing continuum-DSMC methods is its generality: micro elements of any size can be placed at any location (i.e. close to walls or in the bulk of the domain) independently of the continuum grid. The micro resolution can be adjusted to suit each problem, and problems with any degree of spatial scale separation can be simulated. The size

of the micro elements is also able to adapt dynamically depending on the local mean free path of the gas; this will be demonstrated in the following chapters. In addition, the coupling strategy considers the transfer of energy (which has often been neglected in previous methods) and so non-isothermal flows can be dealt with. The flexibility offered by the method makes it applicable for a wide variety of flow problems.

The following two chapters each consider a 1D validation problem, comparing the results from the hybrid method against those obtained from full DSMC simulations of the same problems. The micro Fourier flow problem that will be discussed in Chapter 5 will validate and test the energy coupling strategy. The high-speed micro Couette flow problem considered in Chapter 6 will then validate and assess the performance of the full coupling algorithm. With velocity slip/temperature jump and the accompanying Knudsen layers, both of these 1D flow problems could be considered Type A. However, by assuming that the transport property models are unknown, these problems become Type B — local corrections close to the walls are no longer sufficient and global constitutive corrections are required.

A larger and more realistic engineering problem is considered in Chapter 7: using a microchannel with a high-aspect-ratio cross-section as a representative geometry, the flow of gas through narrow microscale cracks is studied. By assuming periodicity in the streamwise direction, the application of the method is simplified to 2D. This type of high-aspect-ratio geometry can be considered as a specific class of Type C problem and, as mentioned previously, cannot be tackled by DD, HMM, or IMM hybrid methods.

Chapter 5

Micro Fourier flow: a 1D validation problem

The hybrid method is first demonstrated on a simple micro Fourier flow problem; a motionless gas (in this case monatomic argon) is confined between two infinite parallel planar walls that have different temperatures, T_{cold} and T_{hot}. The implementation of the method's coupling algorithm for this particular flow problem will be discussed, and then the results from a number of different test cases will be compared with the results from equivalent full DSMC simulations. Note that, with negligible transport of mass and momentum, the method's coupling is performed via the 1D heat flux only: local temperature fields are imposed on the micro elements, and the resulting global heat flux correction field is supplied to the modified energy equation.

5.1 Computational set-up

Discretising the domain in 1D space, the continuum mesh consists of M_x macro nodes (including a node at each wall). Figure 5.1 shows a schematic of the computational domain; an example micro element arrangement is shown with an element at each wall and one in the bulk. In practice, the appropriate micro element arrangement will depend on the case itself; this will be discussed further in section 5.3.1.

As depicted in Fig. 5.1, the consideration of a 1D problem means that each near-wall element comprises a single sampling region and a single relaxation region, while

each bulk element consists of a single sampling region with a relaxation region on either side. Each sampling region is divided into a number of measurement bins, and each relaxation region is split into a number of control bins. For this flow problem, the measurement and control bins are set to have the same width δx, and the centre of each bin b is set to coincide exactly with a macro node i. The bin width δx is then equal to the spacing between each macro node Δx. Note that this collocation of the bins and macro nodes is not a requirement of the method; it does, however, greatly simplify the transfer of data between the continuum mesh and the micro elements.

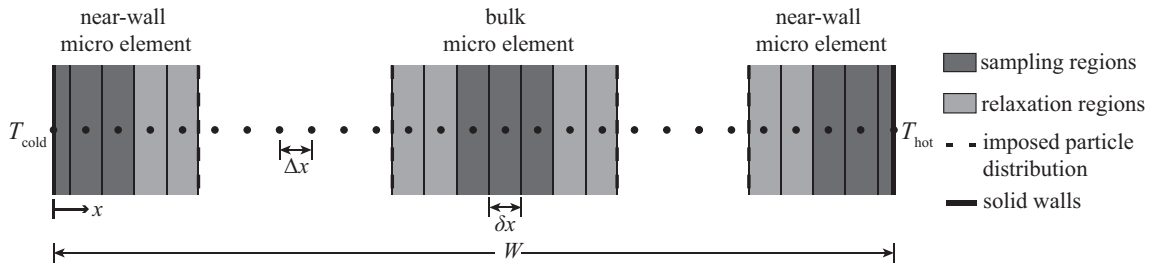

Figure 5.1: Schematic of the computational set-up for a 1D Fourier flow problem.

5.2 Coupling algorithm

To demonstrate the method's ability to deal with missing constitutive information, the variation of the gas thermal conductivity κ with temperature is assumed to be unknown. A reasonable reference value κ_{ref} is adopted, which is constant and independent of the temperature field across the system — the true variation of κ will then be modelled indirectly via the heat flux correction field Φ_x. With this reference conductivity, the modified energy equation (4.7) is reduced to,

$$\kappa_{\text{ref}} \frac{d^2T}{dx^2} - \frac{d\Phi_x}{dx} = 0. \tag{5.1}$$

Rearranging Eq. (4.5), the heat flux correction field Φ_x is obtained from the micro elements according to,

$$\Phi_x = q_x + \kappa_{\text{ref}} \frac{dT}{dx}. \tag{5.2}$$

The coupling algorithm for this validation problem is therefore as follows:

(0) With no heat flux correction, Eq. (5.1) can be approximated using a central finite difference scheme, i.e.

$$\kappa_{\text{ref}} \frac{(T_{i-1} - 2T_i + T_{i+1})}{\Delta x^2} = 0, \tag{5.3}$$

where i=2, 3, ...$(M_x - 1)$. Assuming that there is no temperature jump at the solid walls (i.e. T_1=T_{cold} and T_{M_x}=T_{hot}), solution of this equation produces an initial NSF temperature field T_{NSF}.

(1) All DSMC micro elements are initialised at equilibrium: the temperature and density for initialisation are obtained from the local continuum solution.

Each element is then constrained as follows:

(a) The local continuum temperature field is imposed across each relaxation region by implementing the temperature controller in each control bin.

(b) Maxwellian particle distributions are imposed at the outer boundaries of each relaxation region via a diffuse reflection boundary at the local continuum temperature.

(2) The DSMC algorithm is executed in each micro element. When steady-state conditions are reached, cumulative measurements are performed in each sampling region for some averaging period.

(3) The time-averaged values of the temperature and the heat flux are extracted from each measurement bin. The temperature and the heat flux of the gas at the walls are also extracted from the wall-adjacent bin faces of the near-wall sampling regions.

(4) The heat flux correction in each measurement bin and at each bounding wall is then calculated according to Eq. (5.2). For this calculation, the temperature gradient in each bin (and at each bounding wall) is approximated using a finite difference scheme based on the temperature in adjacent bins (or the adjacent bounding wall), e.g. for a bin b,

$$\Phi_{x_b} = q_{x_b} + \kappa_{\text{ref}} \frac{(T_{b+1} - T_{b-1})}{2\delta x}. \tag{5.4}$$

(5) To obtain Φ_x at every point in the system, linear interpolations are performed between the sampling regions. For this 1D geometry, interpolations of the boundary gas temperature are not required.

(6) Now with the global heat flux correction field Φ_x, Eq. (5.1) can again be approximated using a central finite difference scheme, i.e.

$$\kappa_{\text{ref}} \frac{(T_{i-1} - 2T_i + T_{i+1})}{\Delta x^2} + \frac{(\Phi_{x_{i+1}} - \Phi_{x_{i-1}})}{2\Delta x} = 0, \tag{5.5}$$

where i=2, 3, ...$(M_x - 1)$. Temperature jump is now accounted for: the temperature of the gas at each solid wall is set to be that measured in the corresponding near-wall sampling region during Step (3). Solution of Eq. (5.5) then results in the new, corrected temperature field T_{corr} across the system.

(7) Replacing T_{NSF} with T_{corr}, the process is repeated from Step (1). Iterations continue until T_{corr} converges to within a user-defined tolerance.

5.3 Results

For simplicity, monatomic argon gas is used for all test cases in this chapter: with a molecular mass $m = 66.3\times10^{-27}$ kg and a viscosity exponent $\omega = 0.81$, a reference temperature of 273 K results in a VHS molecular diameter $d_{\text{VHS}} = 4.17\times10^{-10}$ m. The separation W between the two fully diffuse heated walls is maintained at 1 μm for all cases and, to accurately capture the macroscopic variation, $M_x = 201$ macro nodes are used, giving a uniform spacing $\Delta x = 5$ nm $(= \delta x)$.

As was discussed in section 2.1.1, the overall level of non-equilibrium in a system is represented by the global Knudsen number Kn_{gl}. Here, the global mean free path λ_{gl} is calculated from the VHS expression (2.9) based on the average temperature T_{av} and number density n_{av} in the system, while the global characteristic dimension is the separation W, i.e. $Kn_{\text{gl}} = \lambda_{\text{gl}}/W$. By varying Kn_{gl} (i.e. by changing the gas density) and the temperature of the walls, a range of test cases will explore the hybrid method's ability to cope with missing constitutive and boundary information.

To ensure a fair comparison between each hybrid simulation and the corresponding

full DSMC validation simulation, the same cell size and time step is used for the DSMC micro simulations and the full simulation: the cell size is a fraction of the global mean free path λ_{gl}, while the time step δt must be a fraction of the mean collision time t_{mc}. For all cases in this chapter, $\delta t = 1{\times}10^{-12}$ s. An initial start-up run of 3 million time steps enables all simulations to reach steady-state, and a further averaging period of 50 million time steps reduces statistical fluctuations.

The convergence of the hybrid algorithm is monitored using a convergence parameter ζ_T, which quantifies the difference between the temperature solutions in current and previous iterations. Convergence occurs when ζ_T falls below a user-defined tolerance value ζ_{tol_T}, i.e.

$$\zeta_T{}^k = \frac{1}{M_x} \sum_i^{M_x} \left| \frac{T^k_{\text{corr}_i} - T^{k-1}_{\text{corr}_i}}{T^k_{\text{corr}_i}} \right| \leq \zeta_{\text{tol}_T}, \tag{5.6}$$

where i=1, 2, ...M_x and k is the iteration index. The tolerance value ζ_{tol_T} depends on the case itself, with typical values of $\mathcal{O}(10^{-2})$ and below.

A measure of the overall accuracy of the hybrid solution T_{corr} is provided by the mean percentage error $\bar{\epsilon}_T$, comparing against the corresponding full DSMC solution T_{Full}, i.e.

$$\bar{\epsilon}^k_T = \frac{1}{M_x} \sum_i^{M_x} \left[\frac{T_{\text{Full}_i} - T^k_{\text{corr}_i}}{T_{\text{Full}_i}} \times 100\% \right], \tag{5.7}$$

where i=1, 2, ...M_x. The accuracy of the method (and hence the value of this error) depends on how well the micro elements are able to capture the 'true' heat flux correction field across the system. The true heat flux correction is that which can be computed from the full DSMC solution, i.e. $\Phi_{x,\,\text{Full}}$ — this is computed by using the full DSMC temperature field T_{Full} and the full DSMC heat flux field $q_{x,\,\text{Full}}$ in Eq. (5.2).

5.3.1 Micro element configuration

Generally, the number of micro elements and their position and size will determine the accuracy of the method. The appropriate configuration depends on (a) the arrangement of the sampling regions that is needed to provide the required micro resolution, and (b) the size of the relaxation regions that is required to enable full relaxation of the particle state. Typically, particle relaxation occurs over several mean free paths. The required micro resolution is, however, more complex, depending on both the level of

non-equilibrium (globally and locally) and the knowledge of the transport property models.

When modelling this Fourier flow problem, the temperature jump should be captured within the sampling region at each wall. Ideally, each near-wall sampling region should also capture the associated thermal Knudsen layer, which typically extends several mean free paths from a wall surface. Note however that, for high values of Kn_{gl}, the Knudsen layers may be so large that the hybrid approach (over a number of iterations) becomes more expensive than a full DSMC treatment. For lower values of Kn_{gl} where the Knudsen layers are fairly small, micro resolution could be required in the bulk to capture non-equilibrium behaviour caused by strong temperature gradients, or to provide the correct thermal conductivity model.

As mentioned previously, for a given number of micro elements Π, the extents of both the sampling and relaxation regions are able to adapt dynamically depending on the local flow. In this thesis, these extents are set to be some number of local mean free paths λ_l, with λ_l updated at each iteration using Eqs. (2.8) and (2.9). Based on an initial test case, sensitivity studies are now performed in order to explore the appropriate number of local mean free paths for each region.

Setting $Kn_{\mathrm{gl}} = 0.01$, a global mean free path $\lambda_{\mathrm{gl}} = 0.01$ μm corresponds to an average gas number density $n_{\mathrm{av}} = 1.295 \times 10^{26}$ m^{-3}. With an average gas temperature $T_{\mathrm{av}} = 273$ K, a reference value of $\kappa_{\mathrm{ref}} = 0.0164$ W/mK is adopted [121]. The temperature difference ΔT between the heated walls is set to 50 K (i.e. $T_{\mathrm{cold}} = 248$ K and $T_{\mathrm{hot}} = 298$ K), resulting in a temperature gradient of 50×10^6 K/m across the system.

With this value of Kn_{gl} and this temperature gradient, the bulk of the system will be in equilibrium. Although the heat flux correction should be non-zero in the bulk to correct for the error in the assumed conductivity κ_{ref}, the hybrid method should be able to construct an adequate representation of Φ_x by using a linear interpolation between the two near-wall sampling regions. This is due to the fact that, in reality, the thermal conductivity of argon varies approximately linearly with temperature over a range of 200 K to 350 K, and so the true heat flux correction field $\Phi_{x,\,\mathrm{Full}}$ will vary approximately linearly across the bulk of the system. Two near-wall elements are therefore sufficient and no bulk elements are needed, i.e. $\Pi = 2$.

Two studies are performed: in the first, the effect of the relaxation region extent W_{RR} is considered by incrementally increasing its value from $1\lambda_l$ to $7\lambda_l$, while the extent of the sampling regions W_{SR} remains constant at $5\lambda_l$; in the second, the impact of W_{SR} is considered by increasing it from $1\lambda_l$ to $7\lambda_l$, keeping W_{RR} constant at $5\lambda_l$.

With micro elements in this size range, convergence of the hybrid method must occur within 3 to 4 iterations for there to be any computational advantage over a full DSMC simulation. For this test case, a tolerance value of ζ_{tol_T}=0.001 is selected. Figure 5.2(a) shows that, for all W_{RR}, convergence is reached inside 4 iterations. However, as shown in Fig. 5.2(b), convergence does not occur inside 4 iterations when W_{SR} is $1\lambda_l$ or $3\lambda_l$. Larger sampling regions of either $5\lambda_l$ or $7\lambda_l$ are therefore needed, both of which achieve convergence inside 2 to 3 iterations. Note that, for all simulations, the level of convergence fluctuates slightly due to statistical scatter.

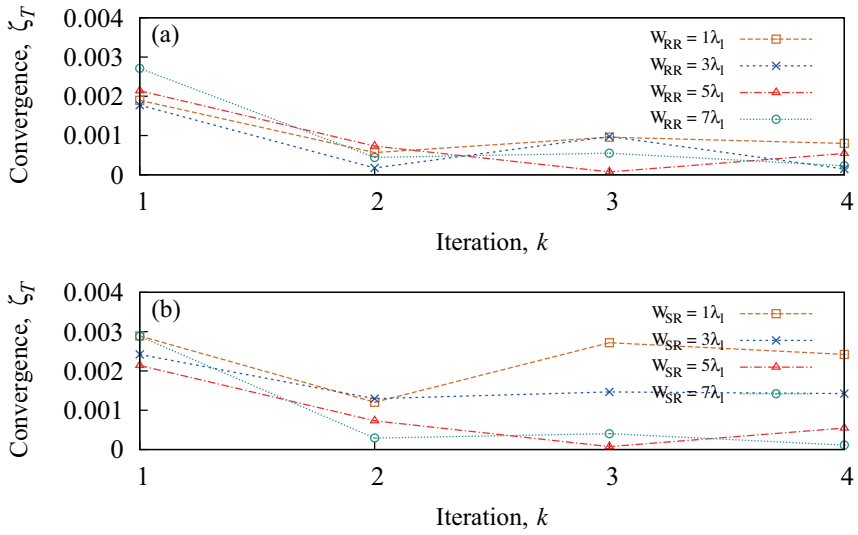

Figure 5.2: Fourier flow: convergence of the hybrid method for (a) W_{SR}=5λ_l and various W_{RR}, and (b) W_{RR}=5λ_l and various W_{SR}, when Kn_{gl}=0.01 and ΔT=50 K.

The final temperature profiles from the hybrid method (i.e. from iteration k=4) are shown in Fig. 5.3; also plotted is the initial NSF and full DSMC temperature fields. The accuracy of the method is, however, more apparent from the mean percentage error $\bar{\epsilon}_T$. Calculated using Eq. (5.7), this error is presented for each iteration in Fig. 5.4; iteration k=0 represents the initial NSF solution. When both the relaxation and sampling region extents are $5\lambda_l$ or $7\lambda_l$, the mean error is reduced to less than 0.1%.

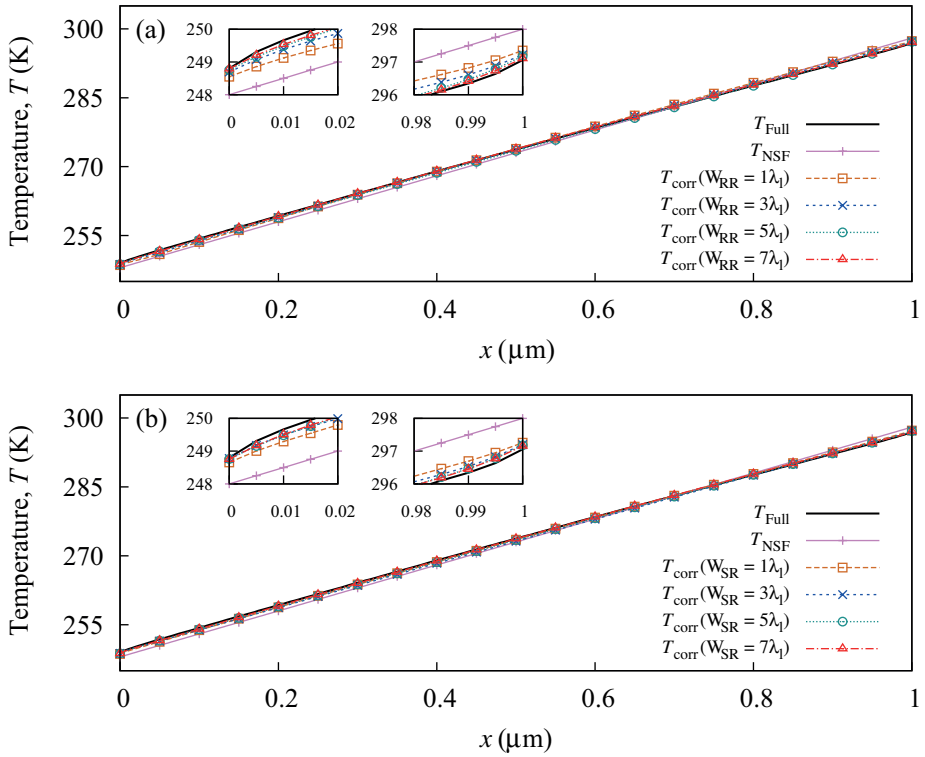

Figure 5.3: Fourier flow: final hybrid temperature solutions T_{corr} for (a) W_{SR}=5λ_l and various W_{RR}, and (b) W_{RR}=5λ_l and various W_{SR}, when Kn_{gl}=0.01 and ΔT=50 K. These are compared with the equivalent full DSMC and NSF temperature solutions. Insets show results at each wall.

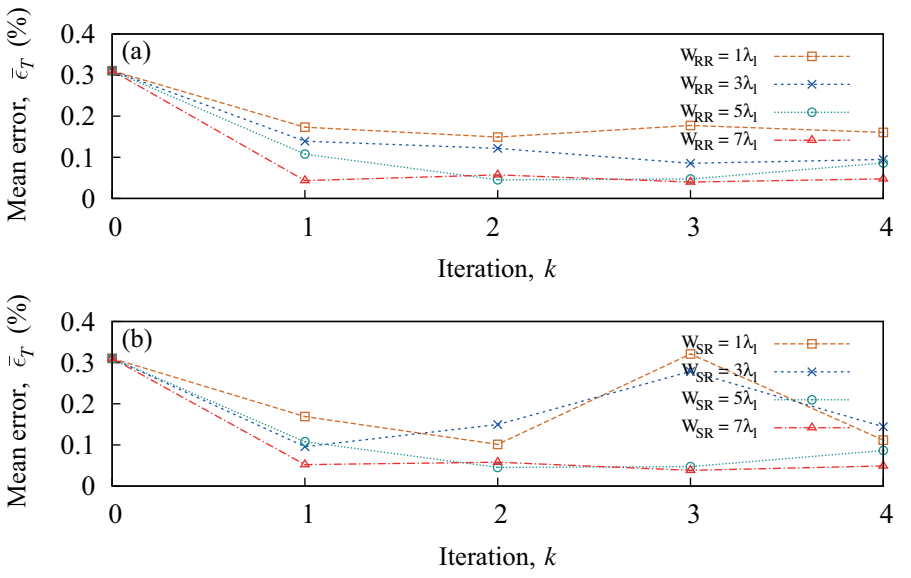

Figure 5.4: Fourier flow: mean error $\bar{\epsilon}_T$ in the hybrid solutions for (a) W_{SR}=5λ_l and various W_{RR}, and (b) W_{RR}=5λ_l and various W_{SR}, when Kn_{gl}=0.01 and ΔT=50 K.

As mentioned previously, the value of $\bar{\epsilon}_T$ depends on the method's ability to capture the true heat flux correction field, i.e. the field computed from the full DSMC solution $\Phi_{x,\,\mathrm{Full}}$. The final heat flux correction fields approximated by the hybrid (at $k = 4$) are shown in Fig. 5.5, along with $\Phi_{x,\,\mathrm{Full}}$.

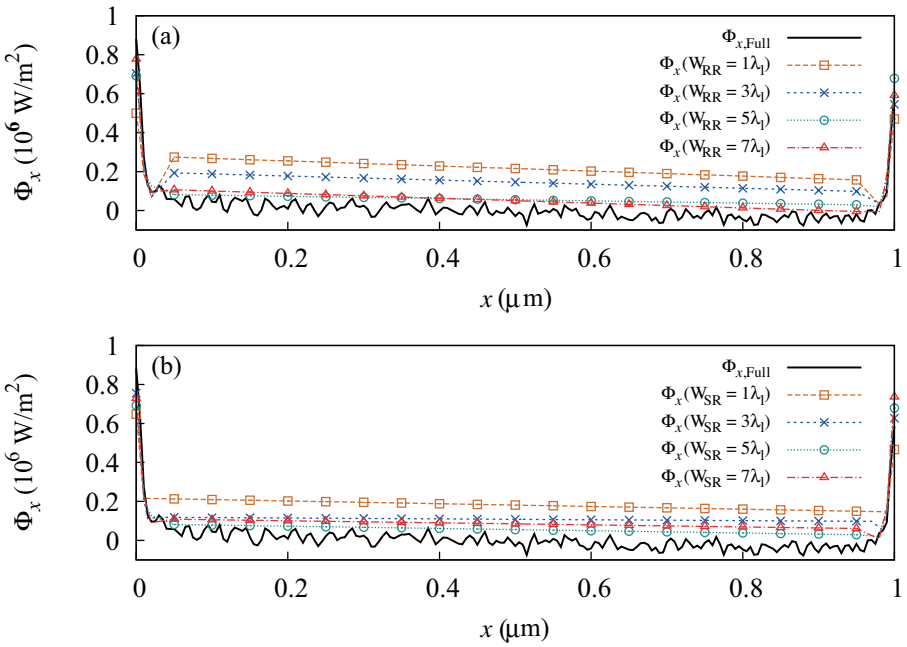

Figure 5.5: Fourier flow: final heat flux correction fields Φ_x for (a) W_{SR}=5λ_l and various W_{RR}, and (b) W_{RR}=5λ_l and various W_{SR}, when Kn_{gl}=0.01 and ΔT=50 K. These are compared with the full DSMC heat flux correction.

Based on these studies, the problems considered in the remainder of this chapter use extents of $5\lambda_l$ for every sampling and relaxation region. Although extents of $7\lambda_l$ would provide slightly higher accuracy, extents of $5\lambda_l$ provide sufficient accuracy while also enabling a greater computational efficiency. Each near-wall element therefore has a total extent of $10\lambda_l$, while each bulk element has a total extent of $15\lambda_l$.

5.3.2 Various rarefaction and temperature conditions

The method's ability to capture the temperature jump and the thermal Knudsen layer is now explored under various rarefaction conditions (study A) and various temperature conditions (study B). In both studies, the wall separation W remains at 1 μm, and the average gas temperature $T_{\mathrm{av}} = 273$ K. By varying the gas density, study A considers three global Knudsen numbers, $Kn_{\mathrm{gl}} = \lambda_{\mathrm{gl}}/W = 0.01, 0.02, 0.03$, while the temperature

difference between the walls remains at ΔT=50 K. Study B considers three temperature differences, $\Delta T = 50, 100, 150$ K, while the density is set to maintain $Kn_{\mathrm{gl}} = 0.01$.

In all six cases, Kn_{gl} and ΔT are small enough that the bulk of the system will be in equilibrium. Also, for argon, the true variation of the thermal conductivity will be approximately linear across these temperature ranges. As in section 5.3.1, two near-wall sampling regions with linear interpolation should therefore be sufficient to capture the heat flux correction field.

With a tolerance value of $\zeta_{\mathrm{tol}_T} = 0.001$, convergence of the hybrid method is reached inside 3 iterations for all six cases; this is shown in Figs. 5.6(a) and 5.6(b) for studies A and B, respectively.

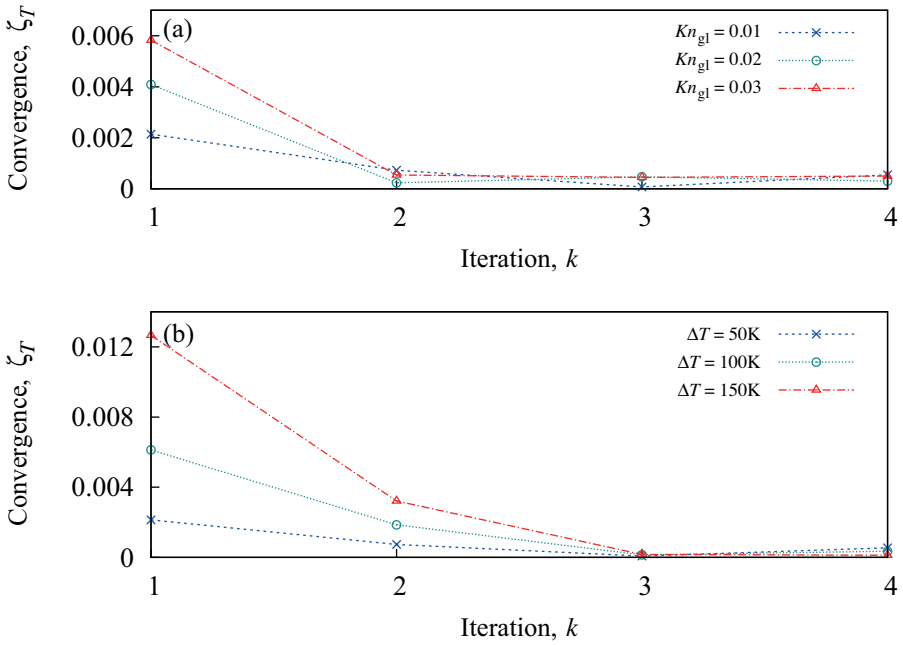

Figure 5.6: Fourier flow: convergence of the hybrid method for (a) various Kn_{gl} (study A), and (b) various ΔT (study B).

The final hybrid temperature solutions (at $k = 4$) from all three cases of study A are presented in Fig. 5.7(a), alongside the corresponding full DSMC solutions. The initial NSF solution is also shown — as this is independent of Kn_{gl}, it is the same for all cases of study A. The final hybrid temperature solutions from all three cases of study B are shown in Fig. 5.7(b); the corresponding NSF and full DSMC solutions are also shown, with the NSF solution different for each value of ΔT.

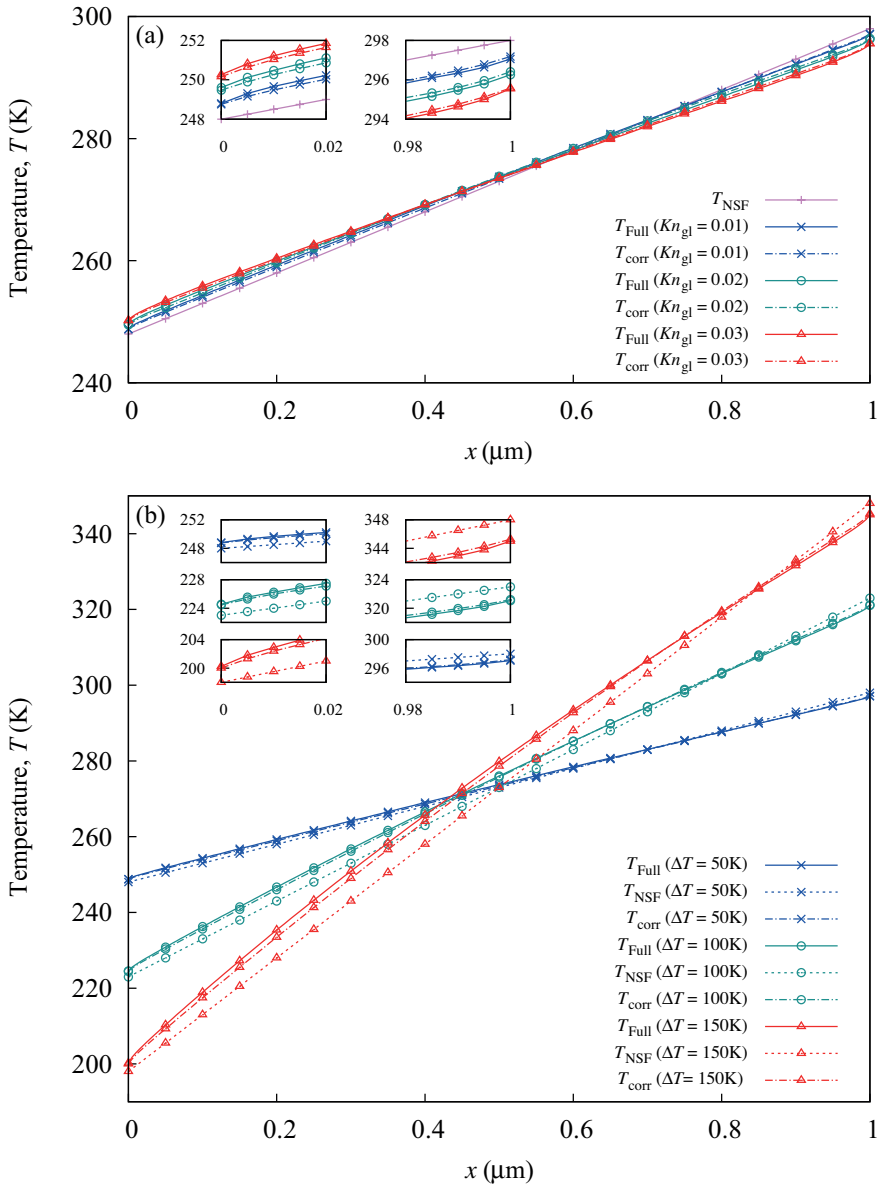

Figure 5.7: Fourier flow: final temperature solutions T_{corr} for (a) various Kn_{gl} (study A), and (b) various ΔT (study B). These are compared with the corresponding NSF and full DSMC solutions. Insets show results at each wall.

Note that the full DSMC temperature solutions from the three cases of study B do not intersect at x=0.5 as the NSF temperature solutions do — this is because the extent of the thermal Knudsen layer at each wall depends on the wall temperature.

Once again, a clearer measure of the hybrid's accuracy is provided by the mean percentage error $\bar{\epsilon}_T$ in the temperature solution, compared with the corresponding full DSMC solution. This error is presented in Fig. 5.8 for each iteration, where iteration

k=0 again represents the initial NSF solution. Figure 5.8(a) shows that, for all three values of Kn_{gl} in study A, the hybrid technique reduces $\bar{\epsilon}_T$ to approximately 0.07%. This is due to the fact that, for all Kn_{gl}, the heat flux correction representation Φ_x is in fairly good agreement with the corresponding full DSMC correction $\Phi_{x,\mathrm{Full}}$, as indicated in Fig. 5.9(a). Figure 5.8(b) shows that the overall accuracy of the hybrid method decreases as ΔT is increased in study B — this is a reflection of the lower quality estimation of the heat flux correction field, as observed in Figure 5.9(b). It should be noted, however, that the hybrid solution provides an improvement over the corresponding NSF solution for all values of ΔT.

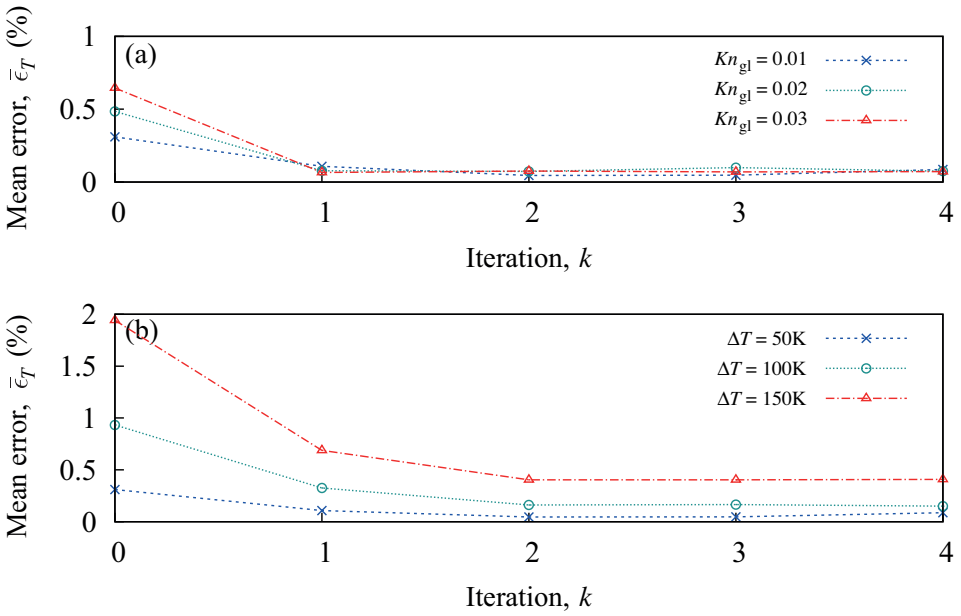

Figure 5.8: Fourier flow: mean error $\bar{\epsilon}_T$ in the hybrid solutions for (a) various Kn_{gl} (study A), and (b) various ΔT (study B).

When modelling higher temperature differences, the accuracy of the method could be improved by increasing the extent of the near-wall sampling regions, or by adding an element in the bulk. This would, however, increase the computational expense. There must be a balance exercised between the accuracy required of, and the computational savings offered by, the hybrid method (particularly for flow problems with low spatial scale separation such as this).

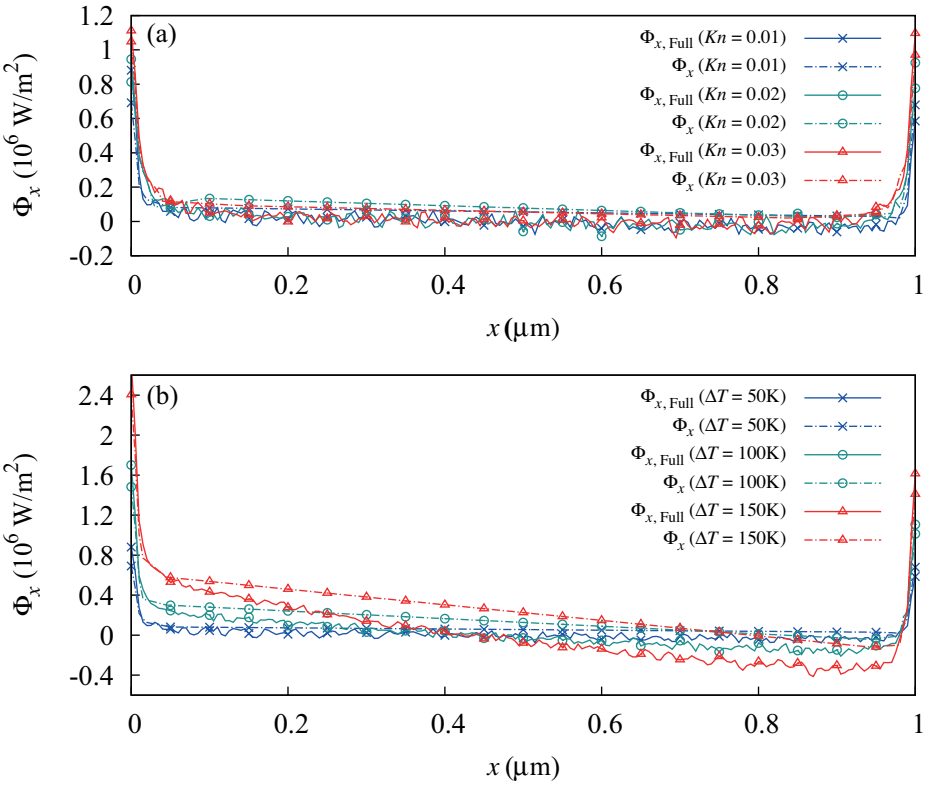

Figure 5.9: Fourier flow: final hybrid heat flux correction fields Φ_x for (a) various Kn_{gl} (study A), and (b) various ΔT (study B), compared with corresponding full DSMC correction fields.

5.3.3 Extreme temperature conditions

For the cases in sections 5.3.1 and 5.3.2, the hybrid method has used only two near-wall micro elements. However, even if the bulk of the system remains close to equilibrium, higher temperature conditions might mean that the true variation of the conductivity with temperature is non-linear. If so, a linear interpolation of the heat flux correction between the two near-wall sampling regions is not likely to produce an accurate solution; the accuracy should be improved by the addition of micro elements in the bulk.

The method's ability to deal with more extreme temperature conditions is now verified on a test case. With $Kn_{\text{gl}} = 0.01$, the average gas temperature T_{av} is increased to 500 K, and so κ_{ref} is increased to 0.027 W/mK [121]. The temperature difference ΔT between the walls is then set equal to 600 K (i.e. $T_{\text{cold}} = 200$ K and $T_{\text{hot}} = 800$ K). Retaining linear interpolations of Φ_x, two micro element configurations are investigated: the first uses only two near-wall elements, i.e. $\Pi = 2$; the second considers the addition of a bulk element in the centre of the system, i.e. $\Pi = 3$.

Setting ζ_{tol_T}=0.01, Fig. 5.10 shows that convergence occurs inside 3 iterations for both configurations. The resulting hybrid temperature fields, along with the initial NSF and the full DSMC temperature fields, are presented in Fig. 5.11.

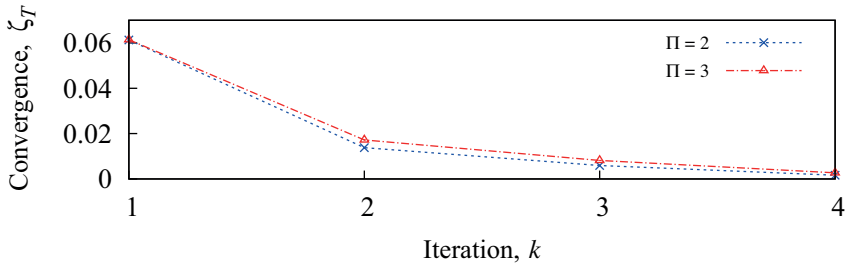

Figure 5.10: Fourier flow: convergence of the hybrid method for Π=2 and Π=3, when Kn_{gl}=0.01 and ΔT=600 K.

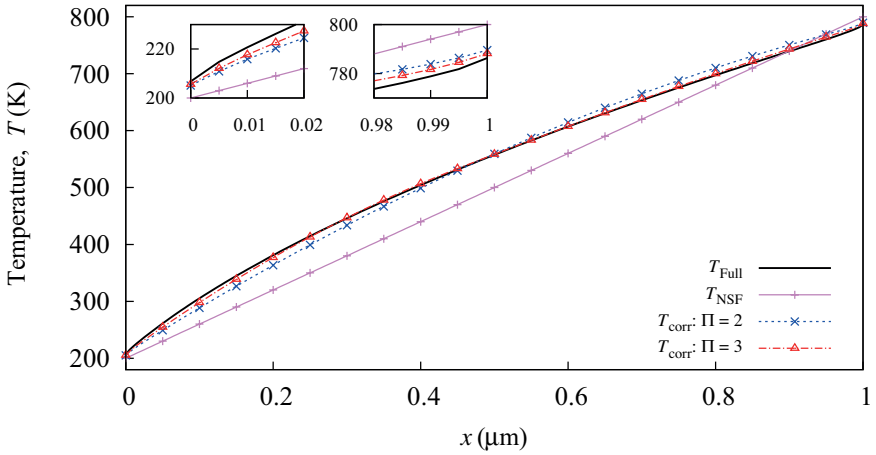

Figure 5.11: Fourier flow: final hybrid temperature solutions for Π=2 and Π=3, alongside the NSF and full DSMC temperature solutions, for Kn_{gl}=0.01 and ΔT=600 K. Insets show results at each wall.

For both configurations, Fig. 5.12 shows the mean percentage error $\bar{\epsilon}_T$ at each iteration: as would be expected, the addition of a bulk element results in a more accurate solution. In fact, in comparison with the initial NSF solution, this configuration provides an order of magnitude reduction in $\bar{\epsilon}_T$. This increased accuracy is the result of the higher quality representation of the heat flux correction, as shown in Fig. 5.13.

The accuracy of the hybrid method could be improved by a further increase in the micro resolution, either by increasing the element extents, or by adding more bulk elements (or both). However, even with only two near-wall micro elements, the method is able to provide a substantial improvement over the NSF solution.

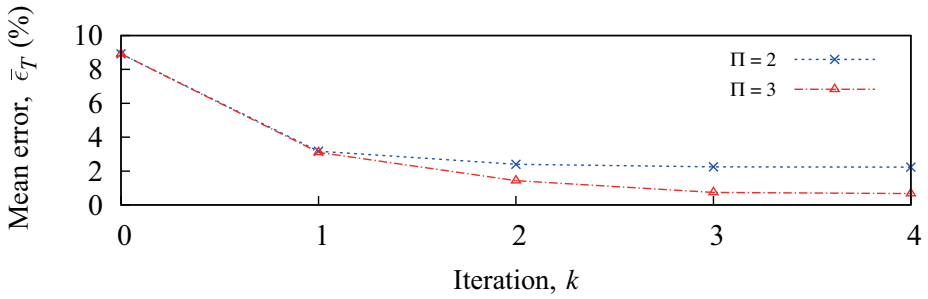

Figure 5.12: Fourier flow: mean error $\bar{\epsilon}_T$ in the hybrid temperature solution with Π=2 and Π=3, when Kn_{gl}=0.01 and ΔT=600 K.

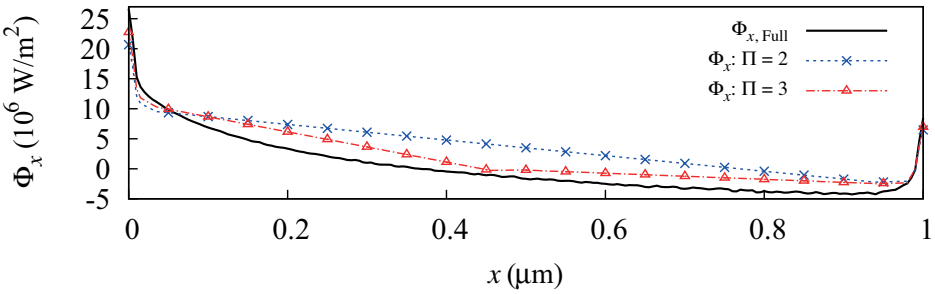

Figure 5.13: Fourier flow: final hybrid heat flux correction fields Φ_x for both Π=2 and Π=3, alongside the full DSMC heat flux correction field, for Kn_{gl}=0.01 and ΔT=600 K.

5.4 Computational savings

In this thesis, computational savings from the hybrid method come from spatial scale separation only — time scale separation has not been exploited, and both the full and micro element DSMC simulations run for the same number of DSMC time steps. However, the dimensions of these Fourier flow cases have been restricted by the need to perform equivalent full DSMC simulations (in order to validate the hybrid method). The thermal Knudsen layers are therefore relatively large, and the micro elements occupy a significant portion of the system. With low spatial scale separation, the computational savings offered by the method are modest.

A measure of the computational speed-up S is given by the ratio of the total processing time required for the full DSMC simulation, to the total processing time required for the hybrid approach. For the hybrid approach, the time taken to compute the continuum solution is negligible compared with the computational time required by the DSMC micro simulations. The total processing time of each approach can therefore

be computed from the total number of DSMC time steps G_{tot} (over both the transient and averaging periods) $\times$ the average clock time per DSMC time step t_{cl}, i.e.

$$S = \frac{G_{\text{tot,Full}} t_{\text{cl,Full}}}{G_{\text{tot,Hyb}} t_{\text{cl,Hyb}}}. \tag{5.8}$$

For all cases in this chapter, $G_{\text{tot,Hyb}} = G_{\text{tot,Full}} = 53{\times}10^6$. The expense of the hybrid method depends on both the micro element configuration and the number of iterations required. As the physical extent of an element can differ in each iteration, the total average clock time per time step for the hybrid approach $t_{\text{cl,Hyb}}$ is calculated as the sum of t_{cl} for all micro elements h=1, 2, ...Π, over all iterations k=1, 2, ...I,

$$t_{\text{cl,Hyb}} = \sum_{k=1}^{I} \left[\sum_{h=1}^{\Pi} t_{\text{cl}}(h,k) \right]. \tag{5.9}$$

Note that, in this thesis, the values of $t_{\text{cl,Full}}$ and $t_{\text{cl,Hyb}}$ that are used to calculate the speed-up S are the average clock times per DSMC time step based on performing both the micro and full DSMC simulations in serial on the ARCHIE-WeSt High Performance Computer; S could therefore be increased by exploiting the scalability of the *dsmcFoamStrath* code in future.

For all cases in section 5.3.2, the hybrid method uses only two near-wall elements and converges inside 3 iterations. With $\Pi = 2$ and $I = 3$, the computational speed-up S for each case is presented in Table 5.1. As the number of particles in a micro element depends on its extent, these speed-ups are determined by the extents of both elements, $W_{\text{left}} = (W_{SR} + W_{RR})_{\text{left}}$ and $W_{\text{right}} = (W_{SR} + W_{RR})_{\text{right}}$, each of which are $10\lambda_l$; the extents presented in Table 5.1 are those at k=3 and subsequent iterations. If Kn_{gl} is 0.01, W_{left} and W_{right} are small enough that a modest speed-up is obtained, i.e. $S > 1$. However, when Kn_{gl} is increased to 0.02 and 0.03, W_{left} and W_{right} become so large that there is no speed-up at all, i.e. $S < 1$.

Convergence is also reached inside $I = 3$ iterations for both micro element configurations considered in section 5.3.3, where $Kn_{\text{gl}} = 0.01$ and $\Delta T = 600$ K. The speed-ups for $\Pi = 2$ and $\Pi = 3$ are shown in Table 5.2. With only near-wall elements, W_{left} and W_{right} are again small enough to provide a modest speed-up; note the change from their original extent of 0.1 μm. Unfortunately, the increase in accuracy that comes from the

addition of the bulk element is also accompanied by an increased computational cost. In fact, with an extent $W_{\text{bulk}} = (W_{SR} + 2W_{RR})_{\text{bulk}}$ corresponding to $15\lambda_l$, the hybrid method provides no savings at all and $S < 1$.

	Kn_{gl}	ΔT (K)	W_{left} (μm)	W_{right} (μm)	I	S
Study A						
	0.01	50	0.1	0.1	3	1.77
	0.02	50	0.19	0.21	3	0.88
	0.03	50	0.29	0.31	3	0.59
Study B						
	0.01	50	0.1	0.1	3	1.77
	0.01	100	0.09	0.11	3	1.73
	0.01	150	0.09	0.11	3	1.77

Table 5.1: Computational speed-ups S for the Fourier flow cases of section 5.3.2.

Π	W_{left} (μm)	W_{bulk} (μm)	W_{right} (μm)	I	S
2	0.07	-	0.12	3	1.71
3	0.07	0.15	0.12	3	0.96

Table 5.2: Computational speed-ups S for the Fourier flow configurations of section 5.3.3.

In summary, several of the simple cases in this chapter demonstrate moderate computational savings from using the hybrid method, while a number of cases are more expensive than the equivalent full DSMC simulation. It is important to note, however, that these Fourier flow cases have been chosen simply to test and validate the energy coupling strategy. Simulations of larger, more realistic, 2D/3D problems are expected to highlight any practical computational advantages of the hybrid method.

5.5 Summary

The new hybrid method has been demonstrated on a micro Fourier flow problem, using 1D energy coupling. The method's ability to compensate for inaccurate boundary and constitutive information has been explored across a range of conditions. For most of the cases considered, the hybrid procedure was found to converge inside only 3 iterations. Generally, good agreement with full DSMC simulations was observed, with the exact level of accuracy depending on the micro element arrangement. The computa-

tional speed-ups reported in this chapter were, at best, modest. The micro element arrangement can, however, be adjusted to attain the desired balance of accuracy with efficiency, and this Fourier flow problem was selected only as a simple validation case. Before tackling 2D/3D problems, the method must first be tested on non-stationary flows with simultaneous momentum and energy coupling.

Chapter 6

High-speed micro Couette flow: a 1D validation problem

Interesting flow problems often involve the transport of momentum *and* heat. The method is therefore used to simulate high-speed micro Couette flow in this chapter; a gas is confined between two infinite parallel planar walls that have the same temperature T_{wall}, but are moving at different planar speeds, u_{lower} and u_{upper}. With high speeds (i.e. Mach number $Ma > 0.2$), viscous heating occurs in the gas. This flow problem will demonstrate the full coupling algorithm of the method: the local continuum velocity, temperature, and density fields are imposed on the micro elements, and the resulting global stress and heat flux corrections are applied to the modified conservation equations. Note, however, that mass conservation is satisfied automatically, and so density coupling is required only in the macro-to-micro direction — no DSMC data is needed to satisfy mass conservation. The results from a number of test cases will be validated against solutions from equivalent full DSMC simulations. Micro Couette flow problems can be representative of the flow in, for example, micro gas bearings, or between computer-disk-reader heads and the rotating hard disks [39].

6.1 Computational set-up

The computational set-up for this Couette flow problem is similar to that outlined for the Fourier flow problem in Chapter 5. The continuum grid takes the form of M_y macro

nodes, as depicted in Fig. 6.1. An example micro element arrangement is shown, while the actual arrangement will depend on the case under consideration. The measurement and control bins have the same height δy and, for convenience, the centre of each bin b is set to coincide exactly with a macro node j. The bin height δy is then equal to the macro node spacing Δy.

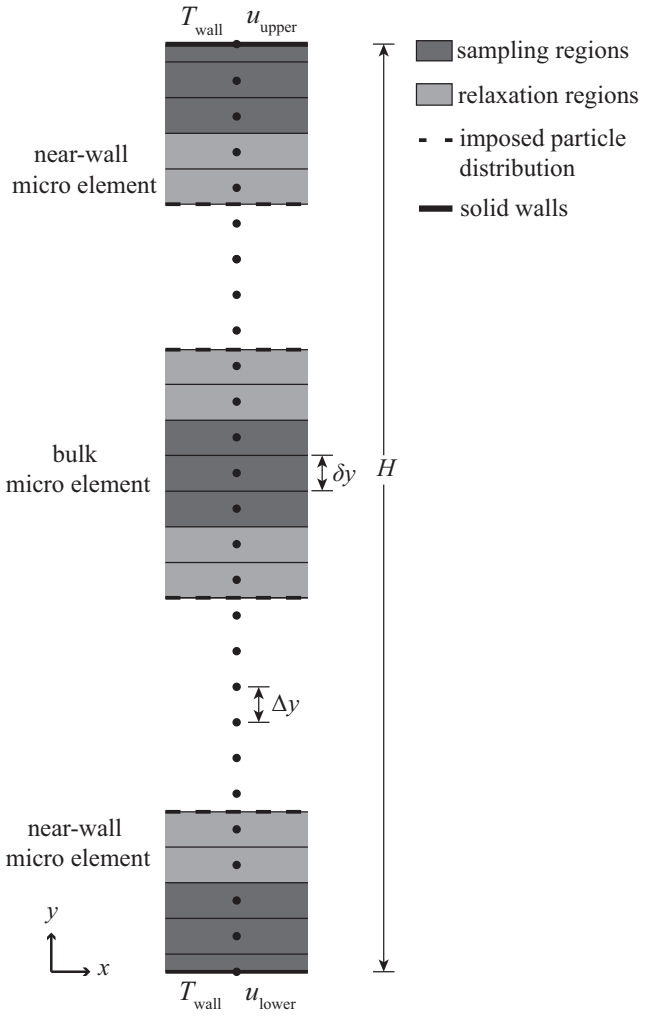

Figure 6.1: Schematic of the computational set-up for a 1D Couette flow problem.

6.2 Coupling algorithm

Reasonable (and constant) reference values are assumed for the gas dynamic viscosity and thermal conductivity: μ_{ref} and κ_{ref}, respectively. The true variation of these trans-

port properties will then be modelled indirectly by the constitutive correction fields. In the streamwise x-direction, the modified momentum conservation equation (4.6) is reduced to,

$$\mu_{\text{ref}}\frac{d^2u}{dy^2} + \frac{d\Omega_{xy}}{dy} = 0, \tag{6.1}$$

where u is the streamwise velocity component. The variation of the pressure transverse to the streamwise direction may be non-zero for high-speed flows, and so the conservation of momentum in the y-direction reads,

$$\frac{dp}{dy} - \frac{d\Omega_{yy}}{dy} = 0. \tag{6.2}$$

With viscous heating in the gas, the conservation of energy must also be considered; the modified energy equation (4.7) is reduced to,

$$\kappa_{\text{ref}}\frac{d^2T}{dy^2} - \frac{d\Phi_y}{dy} + \mu_{\text{ref}}\left(\frac{du}{dy}\right)^2 + \frac{du}{dy}\Omega_{xy} = 0. \tag{6.3}$$

The three constitutive correction fields are computed from the sampling regions of the micro elements by rearranging Eqs. (4.4) and (4.5), i.e.

$$\Omega_{xy} = \tau_{xy} - \mu_{\text{ref}}\frac{du}{dy}, \tag{6.4}$$

$$\Omega_{yy} = \tau_{yy}, \tag{6.5}$$

and

$$\Phi_y = q_y + \kappa_{\text{ref}}\frac{dT}{dy}. \tag{6.6}$$

Therefore, the coupling algorithm for this high-speed Couette problem is as follows:

(0) With no corrections, Eqs. (6.1), (6.2), and (6.3) can be represented by finite difference approximations,

$$\mu_{\text{ref}}\frac{(u_{j-1} - 2u_j + u_{j+1})}{\Delta y^2} = 0, \tag{6.7}$$

$$\frac{(p_j - p_{j-1})}{\Delta y} = 0, \tag{6.8}$$

and

$$\kappa_{\text{ref}}\frac{(T_{j-1}-2T_j+T_{j+1})}{\Delta y^2}+\mu_{\text{ref}}\frac{(u_{j-1}^2-2u_{j+1}u_{j-1}+u_{j+1}^2)}{4\Delta y^2}=0, \tag{6.9}$$

respectively, where j=2, 3, ...(M_y-1). With no velocity slip (i.e. $u_1 = u_{\text{lower}}$ and $u_{M_y} = u_{\text{upper}}$), solving Eq. (6.7) provides the initial NSF velocity field u_{NSF} in the streamwise direction. Unfortunately, there are no obvious initial values for the boundary pressures and so a gauge pressure $p'_1 = 0$ is assumed at the lower wall; solving Eq. (6.8) then produces a gauge pressure field p'_{NSF}. The NSF velocity field u_{NSF} is used to solve Eq. (6.9) — assuming no temperature jump (i.e. $T_1 = T_{M_y} = T_{\text{wall}}$), this provides the NSF temperature field T_{NSF}.

According to the ideal gas law, the actual pressure field p_{NSF} can then be expressed as,

$$p_{\text{NSF}} = p'_{\text{NSF}} + \delta p = \rho_{\text{NSF}} R T_{\text{NSF}}, \tag{6.10}$$

where δp is the difference between the gauge pressure and the actual pressure. Considering the variation of the density field ρ_{NSF} across the channel, i.e.

$$\int \rho_{\text{NSF}} dy = \sum_{j=1}^{M_y} \rho_{\text{NSF}_j} = \rho_{\text{av}} H, \tag{6.11}$$

the pressure difference δp is calculated from,

$$\delta p = \frac{\rho_{\text{av}} H - \frac{1}{R}\sum_{j=1}^{M_y}\left(\frac{p'_{\text{NSF}}}{T_{\text{NSF}}}\right)_j}{\frac{1}{R}\sum_{j=1}^{M_y}\left(\frac{1}{T_{\text{NSF}}}\right)_j}, \tag{6.12}$$

where ρ_{av} is the average density of the gas in the system and H is the separation between the moving walls. The actual pressure field p_{NSF} and the density field ρ_{NSF} are then calculated from Eq. (6.10).

(1) All DSMC micro elements are initialised at equilibrium: the velocity, temperature, and density for initialisation are obtained by averaging over a sub-region of the continuum solution that corresponds to the element.

Each element is then constrained as follows:

(a) The local continuum velocity, temperature, and density fields are imposed over each relaxation region by implementing all three state controllers in each control bin.

(b) Based on the local continuum properties, Maxwellian particle distributions are imposed at the outer boundaries of each relaxation region via a diffuse reflection boundary.

(2) The DSMC algorithm is performed in each micro element for a transient start-up period and a subsequent steady-state averaging period.

(3) The time-averaged macroscopic properties are extracted from the measurement bins of each sampling region. The properties of the gas at the moving walls are also extracted from the wall-adjacent bin faces of both near-wall sampling regions.

(4) The shear stress correction in each measurement bin and at each wall is computed according to Eq. (6.4); the velocity gradient in each bin (and at each bounding wall) is approximated using a finite difference scheme based on the velocities in the adjacent bins (or at the adjacent bounding wall), e.g. for a bin b,

$$\Omega_{xy_b} = \tau_{xy_b} - \mu_{\text{ref}} \frac{(u_{b+1} - u_{b-1})}{2\delta y}. \tag{6.13}$$

From Eq. (6.5), the normal stress correction in each bin b (and at each wall) is simply,

$$\Omega_{yy_b} = \tau_{yy_b}. \tag{6.14}$$

The heat flux correction in each measurement bin and at each wall is computed from Eq. (6.6); again, the temperature gradient is represented by a finite difference approximation, e.g.

$$\Phi_{y_b} = q_{y_b} + \kappa_{\text{ref}} \frac{(T_{b+1} - T_{b-1})}{2\delta y}. \tag{6.15}$$

(5) Interpolations of the correction fields are then performed between the sampling regions to obtain Ω_{xy}, Ω_{yy}, and Φ_y everywhere in the system. Interpolations of the gas boundary information are not required for this 1D problem.

(6) With these global correction fields, Eqs. (6.1), (6.2), and (6.3) can again be represented using a finite difference scheme, i.e.

$$\mu_{\text{ref}}\frac{(u_{j-1}-2u_j+u_{j+1})}{\Delta y^2}+\frac{(\Omega_{xy_{j+1}}-\Omega_{xy_{j-1}})}{2\Delta y}=0, \tag{6.16}$$

$$\frac{(p_j-p_{j-1})}{\Delta y}-\frac{(\Omega_{yy_j}-\Omega_{yy_{j-1}})}{\Delta y}=0, \tag{6.17}$$

and,

$$\begin{aligned}\kappa_{\text{ref}}\frac{(T_{j-1}-2T_j+T_{j+1})}{\Delta y^2}-\frac{(\Phi_{y_{j+1}}-\Phi_{y_{j-1}})}{2\Delta y}\\ +\mu_{\text{ref}}\frac{(u_{j-1}^2-2u_{j+1}u_{j-1}+u_{j+1}^2)}{4\Delta y^2}+\frac{(u_{j+1}-u_{j-1})}{2\Delta y}\Omega_{xy_j}=0,\end{aligned} \tag{6.18}$$

respectively, where j=2, 3, ...(M_y-1). The boundary conditions are now updated to the gas velocities, pressures, and temperatures measured at the bounding walls during Step (3). Solving Eqs. (6.16), (6.17), and (6.18) then provides the new, corrected streamwise velocity field u_{corr}, pressure field p_{corr}, and temperature field T_{corr}, respectively. The new density field ρ_{corr} is also computed from the ideal gas law, $\rho_{\text{corr}} = p_{\text{corr}}/RT_{\text{corr}}$.

(7) Replacing the NSF property fields (i.e. u_{NSF}, T_{NSF}, p_{NSF}, and ρ_{NSF}) with these new, corrected property fields (i.e. u_{corr}, T_{corr}, p_{corr}, and ρ_{corr}), the process is repeated from Step (1). Iterations continue until all fields converge to within user-defined tolerances.

6.3 Results

The working gas is monatomic argon for all test cases in this chapter, with the VHS parameters as stated in section 5.3. The separation H between the walls is 1 μm and $M_y = 201$ macro nodes are used for all cases; this results in a node spacing $\Delta y = 5$ nm ($= \delta y$). Both moving walls are maintained at a temperature $T_{\text{wall}} = 273$ K and are fully diffuse unless stated otherwise. The global Knudsen number $Kn_{\text{gl}} = \lambda_{\text{gl}}/H$ is 0.01 for all cases, i.e. the gas density is set to obtain $\lambda_{\text{gl}} = 0.01$ μm from Eq. (2.9). Although this value of Kn_{gl} is fairly small, viscous heating of the gas will produce local

non-equilibrium effects in the bulk of the flowfield, and the average gas temperature T_{av} will be considerably larger than T_{wall}.

In all cases, both walls move at the same speed but in opposite directions. Two wall speeds are considered to test the coupling strategy of the hybrid method: first, the walls are 'supersonic' with a speed corresponding to a Mach number $Ma \approx 1.1$; secondly, the walls are 'hypersonic' with a speed corresponding to $Ma \approx 5$. Note that these Mach numbers are approximated based on the speed and the temperature of the wall (with no slip/jump), i.e.

$$Ma = \frac{|u_{\text{wall}}|}{\sqrt{\gamma R T_{\text{wall}}}}, \tag{6.19}$$

where $|u_{\text{wall}}| = |u_{\text{lower}}| = |u_{\text{upper}}|$, and γ is the ratio of specific heats.

The same cell dimensions and DSMC time step δt are adopted in both the hybrid micro simulations and the corresponding full DSMC simulation for each test case; for all cases in this chapter, $\delta t = 1{\times}10^{-12}$ s. The start-up and averaging time periods are also the same for both the micro and full DSMC simulations. A start-up run of 3 million time steps enables all DSMC simulations in this chapter to relax to a steady-state. As statistical scatter is lessened at higher flow speeds, the averaging period is set depending on the wall speed: a period of 30 million time steps is used for $Ma \approx 1.1$, while 20 million time steps is adequate for $Ma \approx 5$.

For each property s (i.e. $s = u, T, p$, and ρ), convergence of the hybrid solution occurs when,

$$\zeta_s{}^k = \frac{1}{M_y} \sum_{j}^{M_y} \left| \frac{s^k_{\text{corr}_j} - s^{k-1}_{\text{corr}_j}}{s^k_{\text{corr}_j}} \right| \leq \zeta_{\text{tol}_s}, \tag{6.20}$$

where k is the iteration index, j=1, 2, ...M_y, and ζ_{tol_s} is the property tolerance value that depends on both the property and the case itself; the hybrid algorithm is converged when all properties satisfy their convergence criterion. The accuracy of each hybrid property solution is then quantified by the mean percentage error, i.e.

$$\bar{\epsilon}^k_s = \frac{1}{M_y} \sum_{j}^{M_y} \left[\frac{s_{\text{Full}_j} - s^k_{\text{corr}_j}}{s_{\text{ref}}} \times 100\% \right], \tag{6.21}$$

where s_{ref} is a constant reference value. As the velocity u, temperature T, and density ρ vary considerably across the flowfield, s_{ref} is the range of the full DSMC solution

for these properties, i.e. $s_{\text{ref}} = s_{\text{Full, max}} - s_{\text{Full, min}}$. The pressure p, on the other hand, remains almost constant across the flowfield and so s_{ref} is the average of the full DSMC pressure field, i.e. $p_{\text{ref}} = p_{\text{Full, av}}$. The values of these mean percentage errors will depend on the accuracy of the correction fields Ω_{xy}, Ω_{yy}, and Φ_y constructed by the hybrid method; ideally these should match the full DSMC corrections $\Omega_{xy,\,\text{Full}}$, $\Omega_{yy,\,\text{Full}}$, and $\Phi_{y,\,\text{Full}}$ that are calculated by substituting the full property fields (i.e. u_{Full}, T_{Full}, $\tau_{xy,\,\text{Full}}$, $\tau_{yy,\,\text{Full}}$, and $q_{y,\,\text{Full}}$) into Eqs. (6.4), (6.5), and (6.6), respectively.

6.3.1 Supersonic flow

First, both walls move with a speed corresponding to $Ma \approx 1.1$, i.e. $u_{\text{lower}} = -340$ m/s and $u_{\text{upper}} = 340$ m/s. Assuming no velocity slip or temperature jump, the NSF solution predicts an average gas temperature $T_{\text{av}} = 322.13$ K. Based on this temperature, Eq. (2.9) indicates that an average gas density $\rho_{\text{av}} = 8.98$ kg/m^3 is required to obtain $\lambda_{\text{gl}} = 0.01$ μm. Also based on this temperature, a reference viscosity $\mu_{\text{ref}} = 2.4\times10^{-5}$ kg/ms and reference conductivity $\kappa_{\text{ref}} = 0.019$ W/mK are adopted [121]. A number of micro element configurations are considered for this case; for all, the interpolations of all three correction fields between adjacent sampling regions are simply linear.

An initial configuration is based on the sensitivity study performed in Chapter 5: with only two near-wall elements ($\Pi = 2$), an extent of $5\lambda_l$ is adopted for each sampling and relaxation region, i.e. $H_{SR} = H_{RR} = 5\lambda_l$. As discussed previously, the first iteration of the method assumes $\lambda_l = \lambda_{\text{gl}}$, while subsequent iterations calculate λ_l for each element using Eq. (2.9). Figure 6.2 shows that this initial configuration (represented by the blue dashed line) does not enable the velocity field to converge to within $\zeta_{\text{tol}_u} = 0.02$ inside the 5 iterations performed. The temperature, pressure, and density are, however, able to converge to within $\zeta_{\text{tol}_T} = \zeta_{\text{tol}_p} = \zeta_{\text{tol}_\rho} = 0.002$ inside 3 iterations.

Although it is difficult to observe the accuracy of the final hybrid property fields depicted in Fig. 6.3, the mean percentage errors in these fields shown in Fig. 6.4 confirm that this initial configuration provides only a slight improvement over the NSF solution (represented by results at iteration k=0). This is because of the correction fields constructed by the method, which are presented in Fig. 6.5 for $k = 5$; with this micro element configuration, the approximations of all three correction fields are quite

poor. An increase in the micro resolution should increase the method's accuracy: the addition of bulk elements should improve the capture of the shear stress correction Ω_{xy} and hence the accuracy of u_{corr}; extending the near-wall sampling regions should improve the capture of the heat flux correction Φ_y, increasing the accuracy of T_{corr}.

Three alternative micro element configurations are therefore considered: in the first, there are only two near-wall elements (i.e. $\Pi = 2$); in the second, a bulk element centred at $y = 0.5H$ is added (i.e. $\Pi = 3$); in the third, two bulk elements centered at $y = 0.33H$ and $y = 0.67H$ are added (i.e. $\Pi = 4$). The extents of the near-wall sampling regions are increased to H_{SR}=15λ_l, while each bulk sampling region has an extent of H_{SR}=10λ_l. The extents of the relaxation regions remain at H_{RR}=5λ_l.

With the same tolerance values as the initial test case, Fig. 6.2(a) shows that the velocity convergence is subject to fluctuations when $\Pi = 2$, but occurs inside 3 iterations when $\Pi = 3$ or $\Pi = 4$. Again, the temperature, pressure, and density fields are able to converge inside 3 iterations for all configurations.

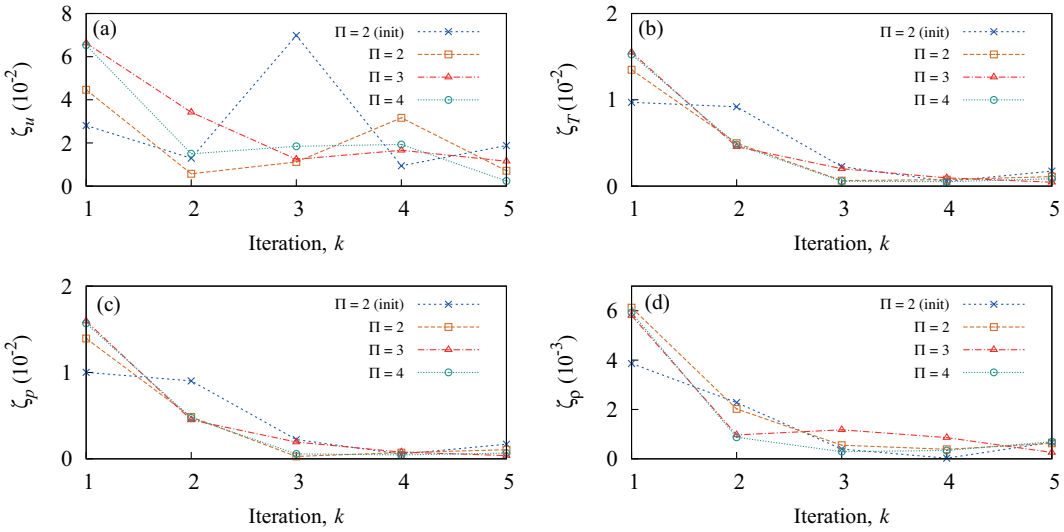

Figure 6.2: Supersonic Couette flow: convergence of the (a) velocity, (b) temperature, (c) pressure, and (d) density solutions for all four of the hybrid configurations implemented.

Compared with the results from the initial test case, the extension of the near-wall sampling regions in the new configurations improves the approximations of all three corrections, as shown in Fig. 6.5. Unfortunately this has come at an increased computational cost, as will be discussed in section 6.4. Adding elements in the bulk offers no real benefit for the capture of Ω_{yy} or Φ_y — it does, however, enable a more

accurate approximation of Ω_{xy}, reducing $\bar{\epsilon}_u$ to below 0.1%. It can be noted that $\Pi = 4$ does not produce significantly better accuracy than $\Pi = 3$; see Fig. 6.4. The use of one bulk element ($\Pi = 3$) can be considered sufficient for this case.

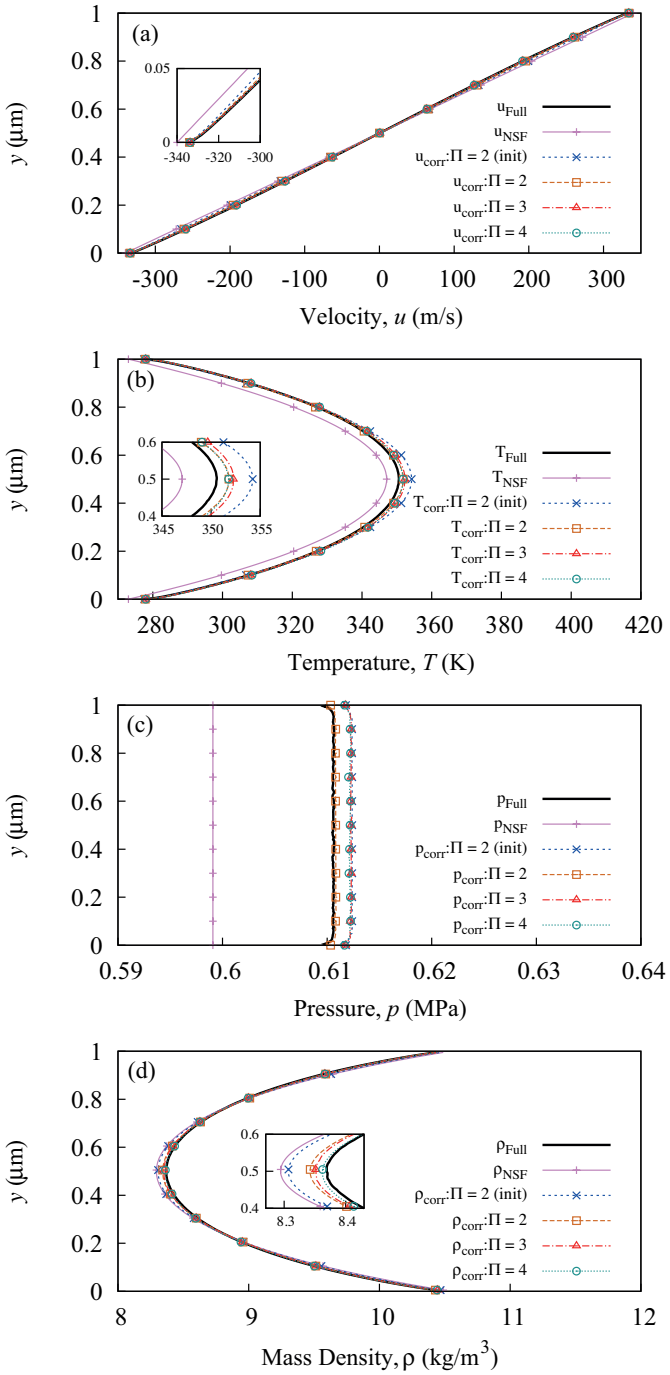

Figure 6.3: Supersonic Couette flow: final hybrid (a) velocity u_{corr} (b) temperature T_{corr} (c) pressure p_{corr} and (d) density ρ_{corr} solutions for all four configurations, alongside the NSF and full DSMC solutions.

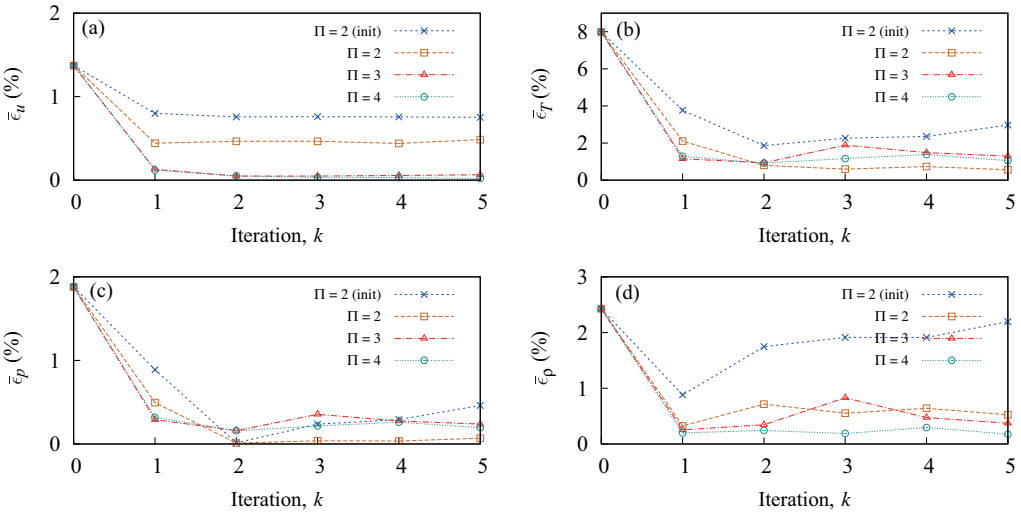

Figure 6.4: Supersonic Couette flow: mean error in the hybrid (a) velocity $\bar{\epsilon}_u$, (b) temperature $\bar{\epsilon}_T$, (c) pressure $\bar{\epsilon}_p$, and (d) density $\bar{\epsilon}_\rho$ solutions for all configurations; the NSF solution is at k=0.

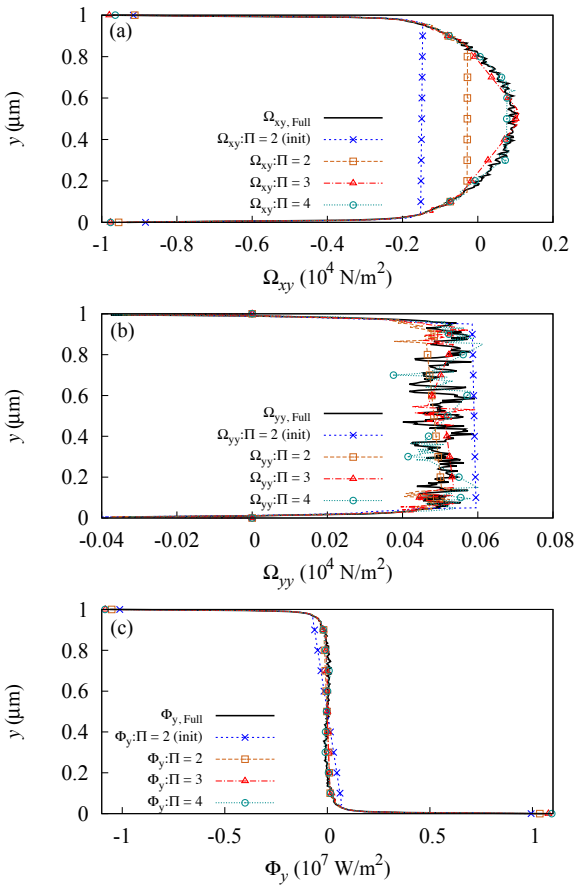

Figure 6.5: Supersonic Couette flow: full DSMC and final (a) shear stress Ω_{xy}, (b) normal stress Ω_{yy}, and (c) heat flux Φ_y correction fields from all four hybrid configurations.

6.3.2 Hypersonic flow

The wall speeds are now increased to $u_{\text{lower}} = -1550$ m/s and $u_{\text{upper}} = 1550$ m/s, corresponding to $Ma \approx 5$. Although hypersonic conditions are unlikely in the types of problems that Couette flow can represent, the purpose of this case is simply to test the hybrid coupling algorithm more rigorously than the previous supersonic flow case.

With no velocity slip or temperature jump, the initial NSF solution predicts an average gas temperature $T_{\text{av}} = 1294$ K. From Eq. (2.9), a gas density $\rho_{\text{av}} = 11.66$ kg/m^3 is then required to obtain $\lambda_{\text{gl}} = 0.01$ μm. Reference values of $\mu_{\text{ref}} = 6.25\times10^{-5}$ kg/ms and $\kappa_{\text{ref}} = 0.049$ W/mK are also assumed [121].

It is expected that the viscous heating in the gas will be significant with such high wall speeds, resulting in highly non-equilibrium flow in the bulk. To test the performance of the hybrid method in such highly non-equilibrium conditions, the same three micro element configurations that were considered for the supersonic flow case are considered once more, i.e. $\Pi = 2$, $\Pi = 3$, and $\Pi = 4$; the near-wall and bulk sampling regions have extents of $15\lambda_l$ and $10\lambda_l$, respectively, while the relaxation regions have extents of $5\lambda_l$. Once again, linear interpolations are performed between the sampling regions to approximate the global corrections.

With tolerance values of $\zeta_{\text{tol}_u} = 0.05$ and $\zeta_{\text{tol}_T} = \zeta_{\text{tol}_p} = \zeta_{\text{tol}_\rho} = 0.002$, Fig. 6.6 shows that all hybrid property fields converge inside 3 iterations, for all three configurations.

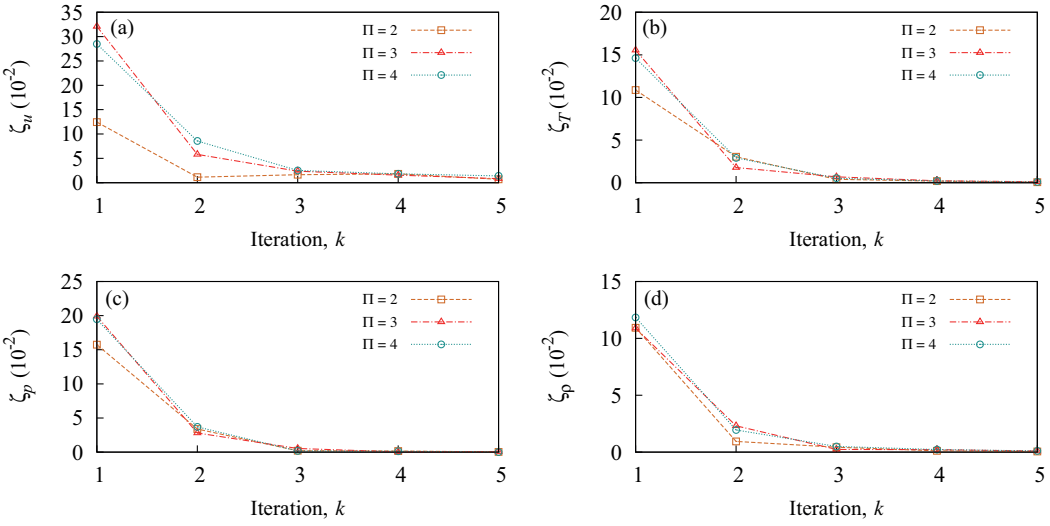

Figure 6.6: Hypersonic Couette flow: convergence of the (a) velocity, (b) temperature, (c) pressure, and (d) density solutions for all three of the hybrid configurations implemented.

The final property fields at iteration $k = 5$ are shown in Fig. 6.7, while the mean percentage error in the property fields at each iteration is presented in Fig. 6.8. The final hybrid correction fields at $k = 5$ are also shown in Fig. 6.9.

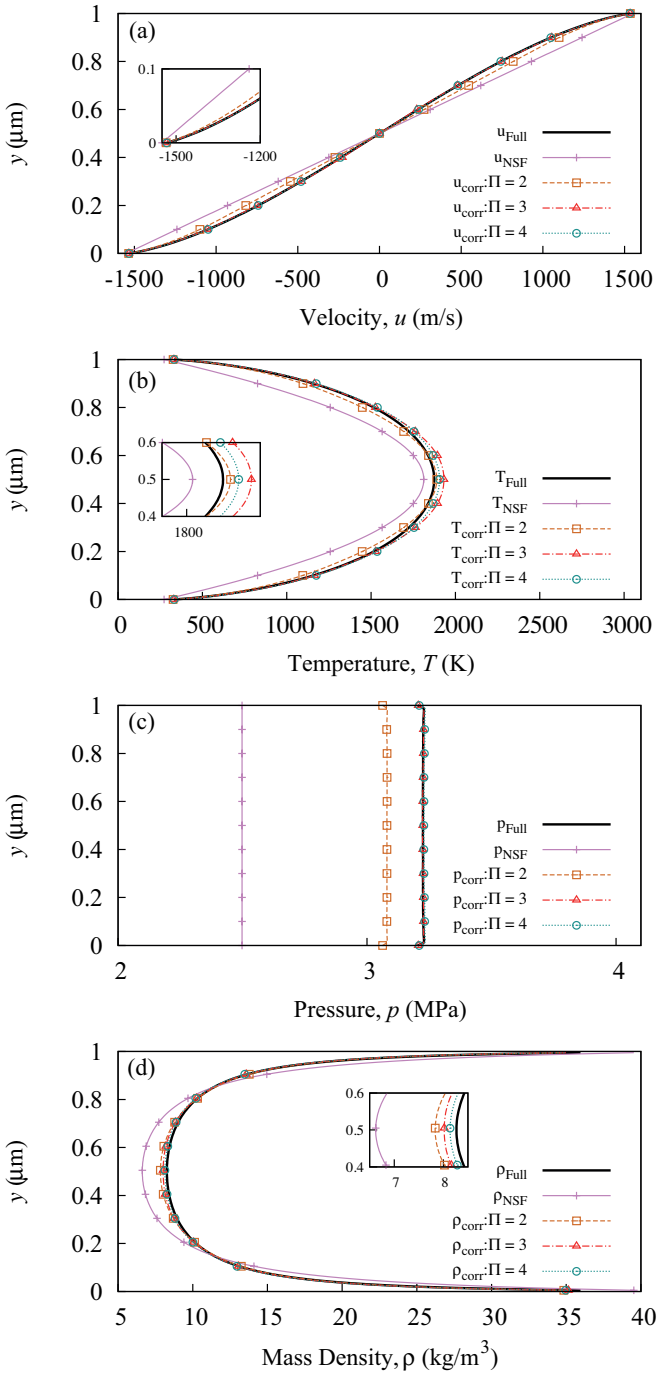

Figure 6.7: Hypersonic Couette flow: final hybrid (a) velocity u_{corr} (b) temperature T_{corr} (c) pressure p_{corr} and (d) density ρ_{corr} solutions from all configurations, alongside the NSF and full DSMC solutions.

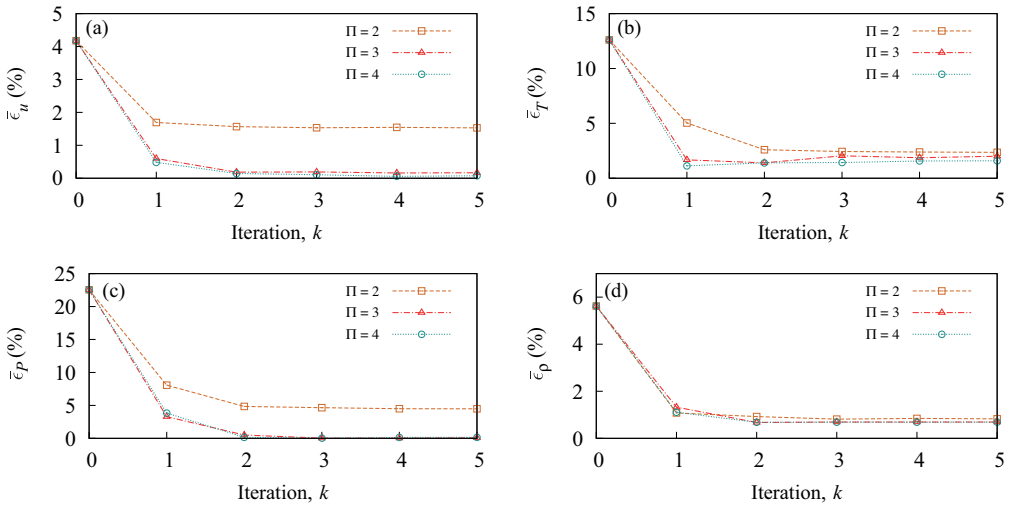

Figure 6.8: Hypersonic Couette flow: mean error in the hybrid (a) velocity $\bar{\epsilon}_u$, (b) temperature $\bar{\epsilon}_T$, (c) pressure $\bar{\epsilon}_p$, and (d) density $\bar{\epsilon}_\rho$ solutions, for all configurations; the NSF solution is at k=0.

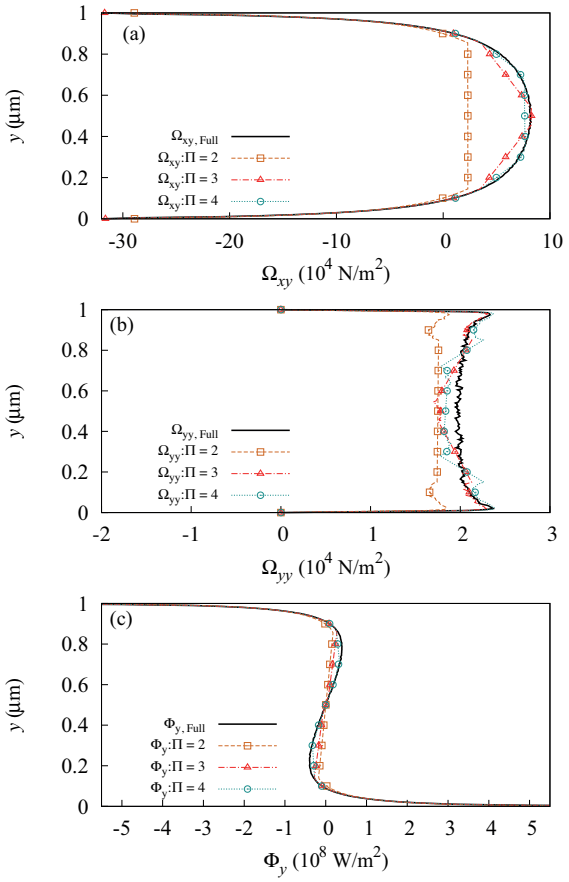

Figure 6.9: Hypersonic Couette flow: full DSMC and final (a) shear stress Ω_{xy}, (b) normal stress Ω_{yy}, and (c) heat flux Φ_y correction fields from all three hybrid configurations.

As would be expected, the higher degree of non-equilibrium in this hypersonic flow case means that (when comparing against the corresponding full DSMC solutions) the NSF prediction of the property fields is considerably less accurate than the NSF prediction for the previous supersonic case, i.e. compare $k = 0$ in Fig. 6.4 and Fig. 6.8. When simulating this hypersonic case with only two near-wall elements, the mean errors in the final velocity and pressure fields produced from the method are $\bar{\epsilon}_u = 1.53\%$ and $\bar{\epsilon}_p = 4.47\%$, respectively. By adding elements in the bulk, improved approximations of Ω_{xy} and Ω_{yy} reduce these errors down to approximately 0.1%, for both $\Pi = 3$ and $\Pi = 4$. The largest difference in accuracy between the two bulk element configurations lies in the temperature solution, largely due to the capture of Φ_y: one bulk element results in $\bar{\epsilon}_T = 2.42\%$, while two bulk elements gives $\bar{\epsilon}_T = 1.59\%$. This increase in accuracy might not be considered to be worth the increase in the computational cost; the cost of all three configurations will be discussed in section 6.4.

Alternative interpolation approaches

So far, all of the hybrid simulations performed in this chapter (and the previous chapter) have approximated the global corrections across the entire flowfield by using simple linear interpolations between the local correction fields of adjacent sampling regions. Based on this hypersonic test case and considering both bulk element configurations (i.e. $\Pi = 3$ and $\Pi = 4$), three alternative approaches are now investigated.

The first is cubic interpolations: essentially, a third degree polynomial is generated to fill the gap between adjacent sampling region corrections. Instead of interpolations, the second alternative fits a global curve across the entire flowfield based on the local corrections in all of the sampling regions: a polynomial curve of order n_p is generated using the MATLAB 'polyfit' function [122]. The third alternative is also a global curve-fitting approach: B-splines are fit over the entire flowfield using the MATLAB 'splinefit' function [123], where the data is split into pieces and splines of order n_s are applied to each piece; here, the correction data is split into 20 uniform pieces (i.e. each piece has the same number of data points).

Before performing any hybrid simulations, a fair comparison of these alternative approaches is made by performing tests on the 'ideal' sampling region corrections, i.e.

interpolation/fitting is carried out based on the sections of the full DSMC correction fields that correspond to each sampling region. Along with observing the resulting global correction fields, the accuracy of each approach is measured by calculating the mean percentage error in comparison with the actual correction fields computed from the full DSMC solution, i.e. for $\Gamma = \Omega_{xy}, \Omega_{yy}$ and Φ_y,

$$\bar{\epsilon}_\Gamma = \frac{1}{M_y} \sum_j^{M_y} \left[\frac{\Gamma_{\text{Full}_j} - \Gamma_j}{(\Gamma_{\text{Full, max}} - \Gamma_{\text{Full, min}})} \times 100\% \right]. \tag{6.22}$$

As all three corrections are known to vary considerably across the system, the denominator for each is the range of the full DSMC correction field.

The values of $\bar{\epsilon}_{\Omega_{xy}}$, $\bar{\epsilon}_{\Omega_{yy}}$, and $\bar{\epsilon}_{\Phi_y}$ that result from applying linear interpolations and these three alternatives approaches to the ideal sampling region data are given in Table 6.1, for both Π=3 and Π=4. Linear interpolations are fairly accurate for both configurations, with the errors in all three correction fields below 1.5% — this approach therefore provides a benchmark, against which the other three approaches are measured. Note that only selected values of n_p and n_s are shown: these highlight the spread of accuracy; higher values are typically less accurate. As would be expected, two bulk elements ($\Pi = 4$) are more accurate than only one ($\Pi = 3$) for all four approaches.

		Π=3			Π=4		
	n_p or n_s	$\bar{\epsilon}_{\Omega_{xy}}$	$\bar{\epsilon}_{\Omega_{yy}}$	$\bar{\epsilon}_{\Phi_y}$ (%)	$\bar{\epsilon}_{\Omega_{xy}}$	$\bar{\epsilon}_{\Omega_{yy}}$	$\bar{\epsilon}_{\Phi_y}$ (%)
Linear	-	1.34	0.61	0.38	0.37	0.59	0.10
Cubic	-	0.68	0.81	0.19	0.26	0.62	0.06
Polyfit	3	6.43	2.55	4.13	6.86	2.54	2.01
	6	7.04	8.31	3.27	3.96	4.96	1.37
	9	11.13	20.50	3.05	2.09	4.27	0.77
	12	5.32	18.58	6.83	2.16	6.64	0.77
	15	13.07	45.84	2.55	1.62	6.26	0.50
	18	27.24	65.51	11.05	2.26	8.95	0.71
Splinefit	3	0.53	1.60	0.34	0.23	1.59	0.27
	6	1.45	2.62	0.47	0.75	1.34	0.13
	9	6.12	14.83	1.11	0.63	2.89	0.13
	12	15.05	24.88	3.82	1.18	5.77	0.22
	15	18.74	30.59	4.60	1.07	6.06	0.18
	18	59.68	248.3	36.62	1.20	26.60	1.40

Table 6.1: Mean errors in the correction fields produced from all interpolation/fitting approaches.

Although it is difficult to observe from the resulting global correction fields shown in Fig. 6.10, the error values in Table 6.1 indicate that, compared with linear interpolations, cubic interpolations provide a slight increase in the accuracy of Ω_{xy} and Φ_y and a slight decrease in the accuracy of Ω_{yy}.

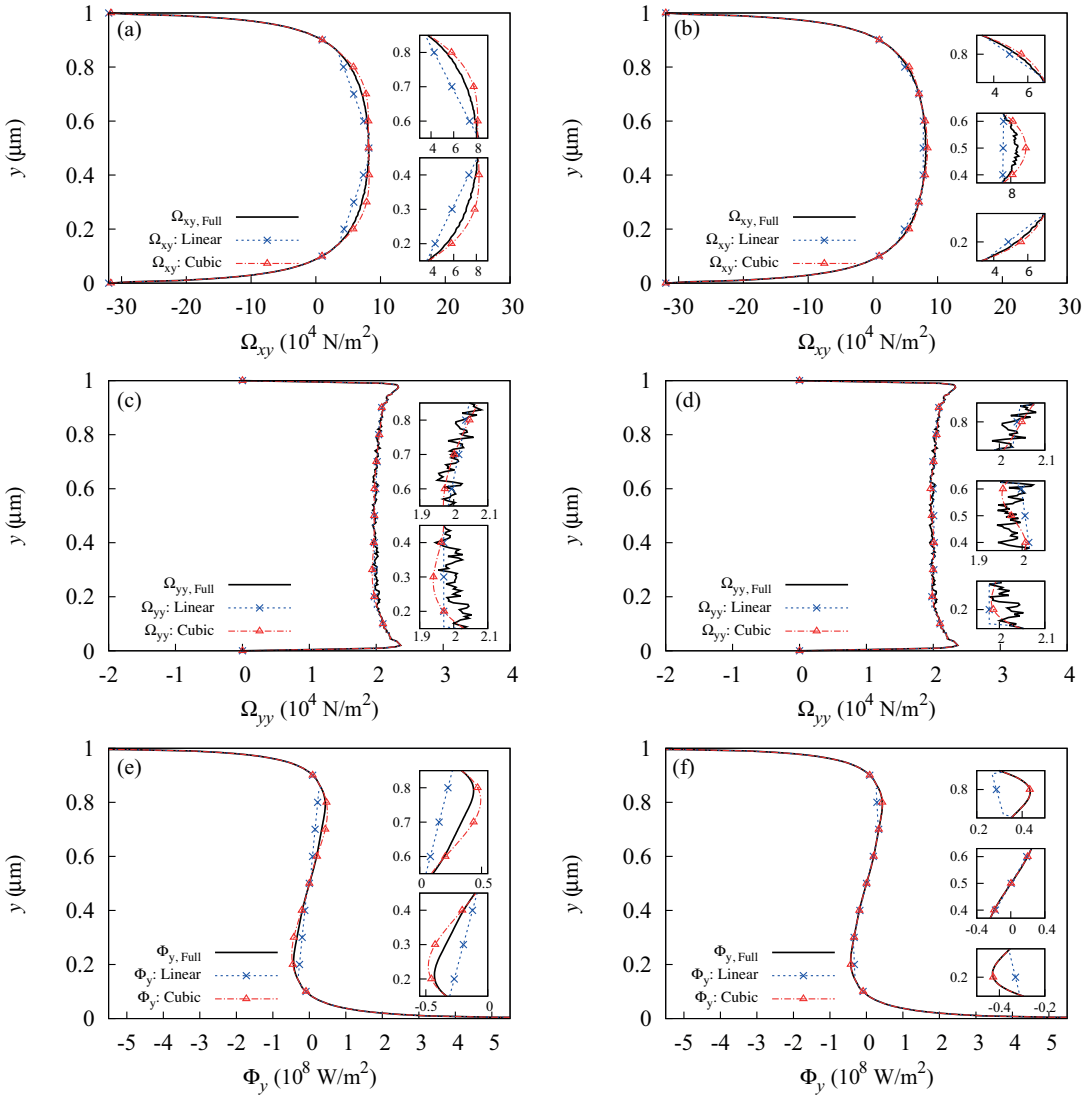

Figure 6.10: Global correction fields with linear and cubic interpolations: Ω_{xy} for (a) $\Pi = 3$ and (b) $\Pi = 4$; Ω_{yy} for (c) $\Pi = 3$ and (d) $\Pi = 4$; and Φ_y for (e) $\Pi = 3$ and (f) $\Pi = 4$. Full DSMC global corrections are also shown. Insets show detailed views of interpolations.

For both $\Pi = 3$ and $\Pi = 4$, the fitting of polynomial and B-spline curves is increasingly erratic with increasing values of n_p and n_s; see Figs. 6.11 and 6.12, respectively. The mean error values in Table 6.1 confirm that, with the exception of $n_s = 3$, the polynomial and spline fitting approaches are less accurate than linear interpolations.

Note that, although the normal stress correction curves Ω_{yy} produced for $n_p = 3$ and $n_p = 6$ appear to match well to the full DSMC correction $\Omega_{yy,\text{Full}}$ in Fig. 6.11, these polynomial curves are not able to capture the sharp change close to the bounding walls. Table 6.1 also shows that spline fitting is typically less accurate than polynomial fitting for $\Pi = 3$, but more accurate than polynomial fitting for $\Pi = 4$. Similar to the cubic interpolations, fitting cubic splines with $n_s = 3$ across the flowfield provides an increase in the accuracy of Ω_{xy} and Φ_y and a decrease in the accuracy of Ω_{yy} when compared with linear interpolations.

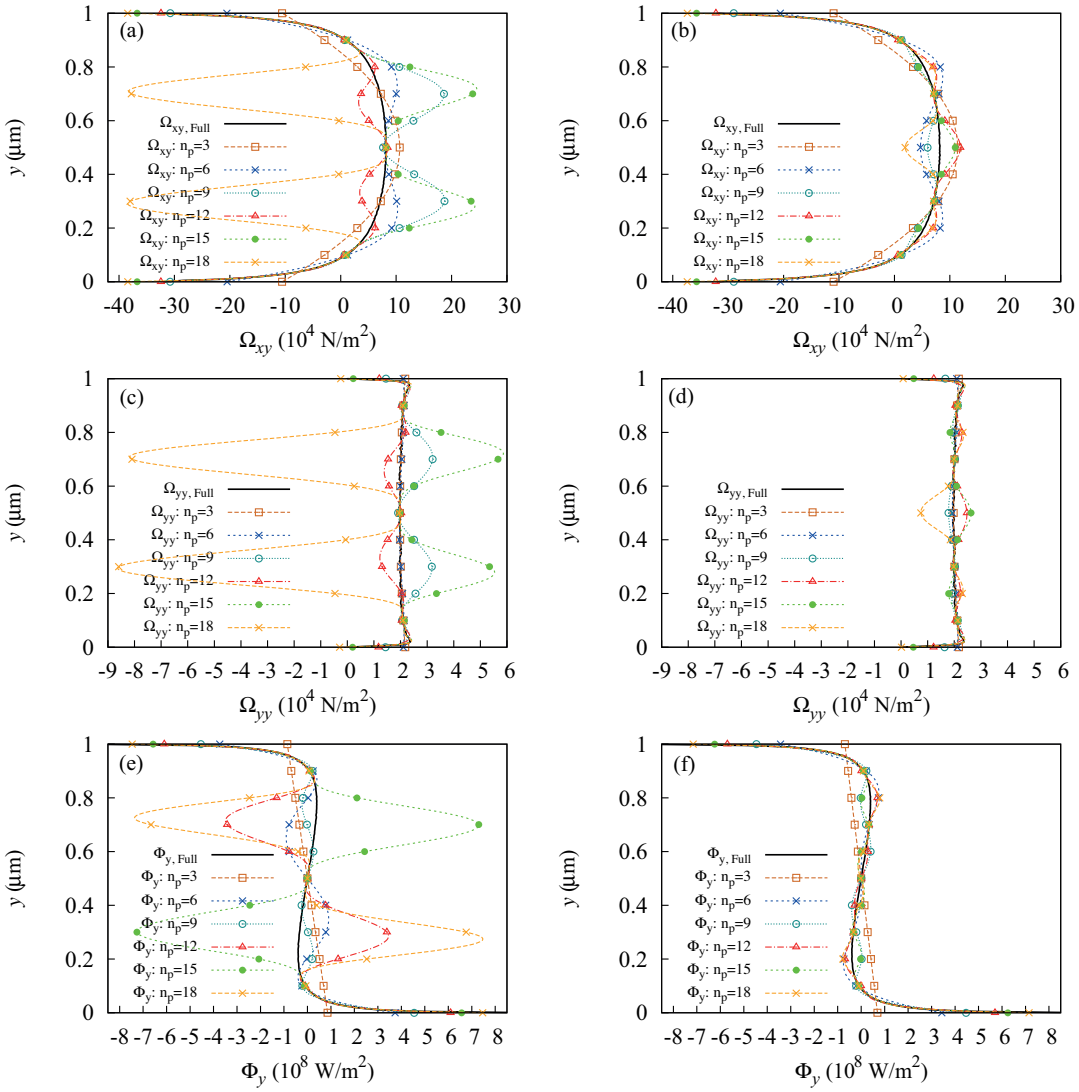

Figure 6.11: Global correction fields with polynomial fitting: Ω_{xy} for (a) $\Pi = 3$ and (b) $\Pi = 4$; Ω_{yy} for (c) $\Pi = 3$ and (d) $\Pi = 4$; and Φ_y for (e) $\Pi = 3$ and (f) $\Pi = 4$. Full DSMC global corrections are also shown.

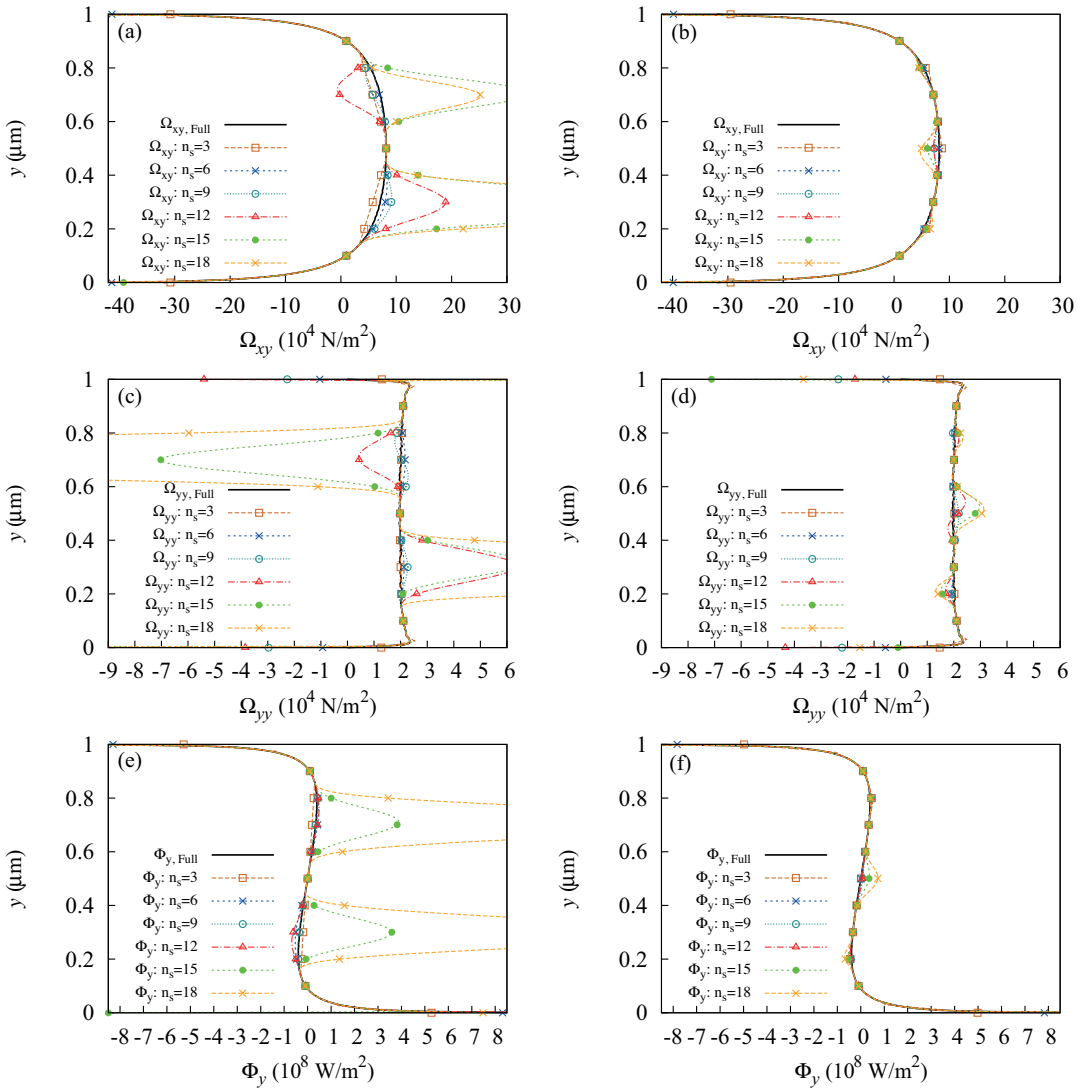

Figure 6.12: Global correction fields with spline fitting: Ω_{xy} for (a) $\Pi = 3$ and (b) $\Pi = 4$; Ω_{yy} for (c) $\Pi = 3$ and (d) $\Pi = 4$; and Φ_y for (e) $\Pi = 3$ and (f) $\Pi = 4$. Full DSMC global corrections are also shown.

Therefore, based on this testing, cubic interpolations and cubic spline fitting (i.e. with $n_s = 3$) were implemented in hybrid simulations of this hypersonic test case; note that these alternative approaches were used only to obtain the global Ω_{xy} and Φ_y fields, with linear interpolations performed to obtain the global Ω_{yy} field.

When comparing the mean percentage errors in the property fields that resulted from these hybrid simulations, cubic interpolations were found to have approximately the same accuracy as linear interpolations, as depicted in Fig. 6.13 for $\Pi = 4$. However, the response of the cubic spline fitting approach to the real sampling region data was

generally unpredictable; see Fig. 6.14 for $\Pi = 4$. Unfortunately, this resulted in the iterative algorithm of the hybrid method becoming unstable, with the property fields (particularly T_{corr}) diverging from the full DSMC solution as iterations continued.

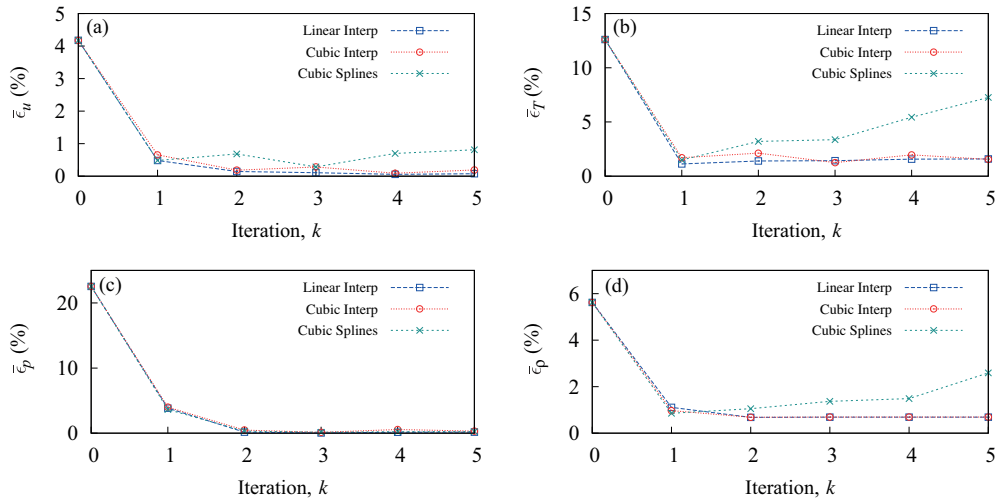

Figure 6.13: Hypersonic Couette flow: mean errors in the hybrid (a) velocity, (b) temperature, (c) pressure, and (d) density solutions with different interpolation/fitting approaches and Π=4.

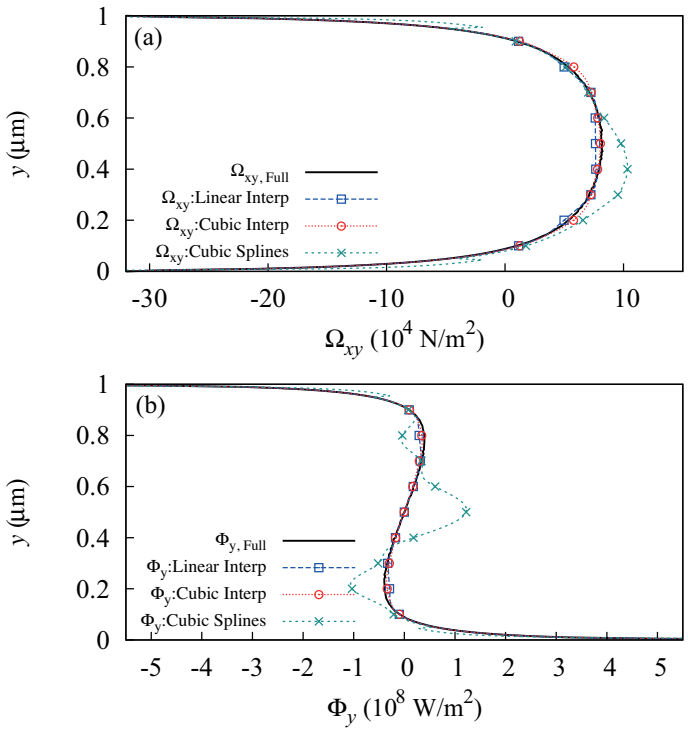

Figure 6.14: Hypersonic Couette flow: full DSMC and final hybrid (a) shear stress Ω_{xy}, and (b) heat flux Φ_y correction fields, with different interpolation/fitting approaches and Π=4.

With no particular advantage from the use of cubic interpolations, linear interpolations are adequate for this case. In fact, for any given problem, simple linear interpolations will provide the most reliable approach to approximating the global correction fields and, with sufficient micro resolution, could produce a fairly accurate solution.

Mixed specular-diffuse walls

The conventional NSF continuum-fluid description with no slip/jump does not consider specular reflection — the bounding walls are essentially assumed to be fully diffuse. The near-wall micro elements of this hybrid method should, however, be able to capture the effects of specular reflection; this is now demonstrated by setting the moving walls of this hypersonic test case to have a diffuse fraction α of 0.5.

To compare the method's performance for this mixed specular-diffuse case with that for the fully diffuse case, the micro element configuration that was the most accurate for the fully diffuse case is adopted, i.e. two near-wall elements and two bulk elements, $\Pi = 4$. Once again, the extents of the near-wall and bulk sampling regions are $15\lambda_l$ and $10\lambda_l$, respectively, and the relaxation region extents are $5\lambda_l$. Following the discussion in the previous sub-section, linear interpolations are used to approximate the global corrections. Figure 6.15 shows the convergence of the hybrid simulations for both types of wall boundaries, with the final property fields presented in Fig. 6.16.

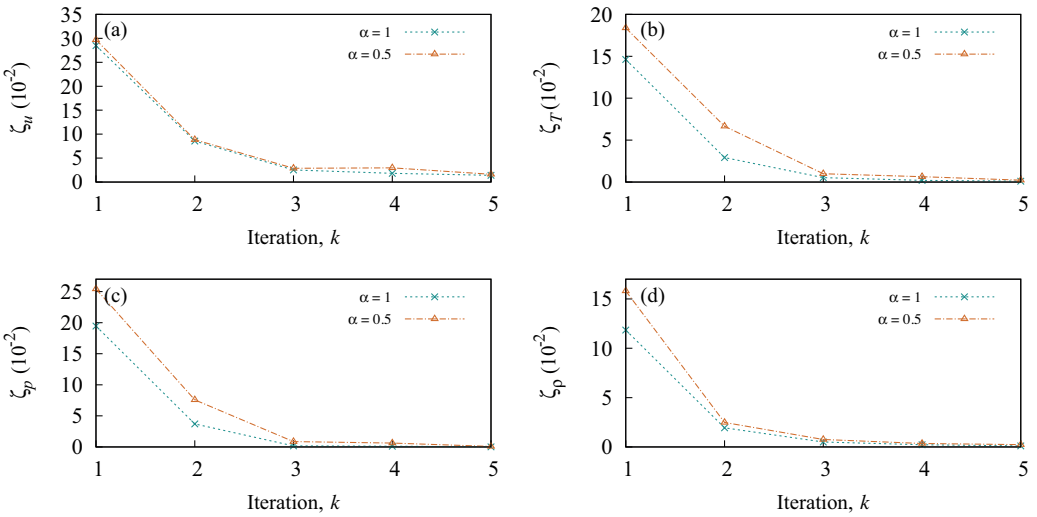

Figure 6.15: Hypersonic Couette flow: convergence of the hybrid (a) velocity, (b) temperature, (c) pressure, and (d) density solutions for fully diffuse (α=1) and mixed specular-diffuse (α=0.5) walls.

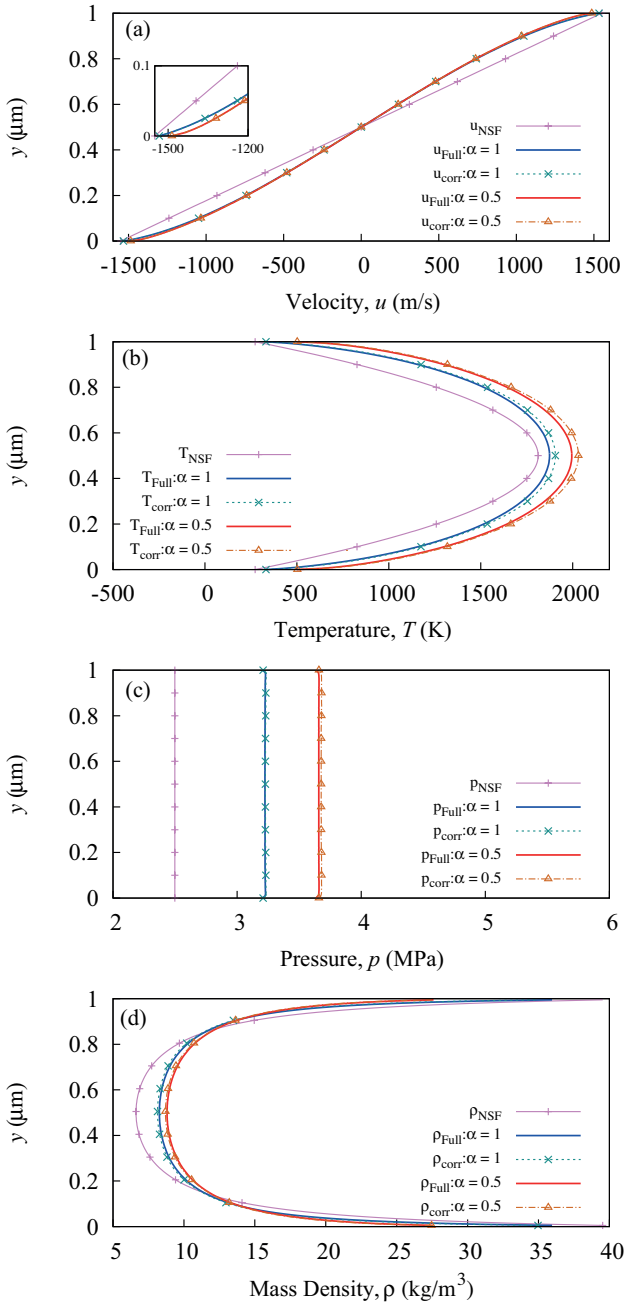

Figure 6.16: Hypersonic Couette flow: full DSMC and final hybrid (a) velocity (b) temperature (c) pressure and (d) density solutions for fully diffuse (α=1) and mixed specular-diffuse (α=0.5) walls; note that the NSF solution is the same for both wall types.

Once again, values of $\zeta_{\text{tol}_u} = 0.05$ and $\zeta_{\text{tol}_T} = \zeta_{\text{tol}_p} = \zeta_{\text{tol}_\rho} = 0.002$ are selected, and so the method converges inside 3 iterations for both cases. Compared with the corresponding full DSMC solutions, the mean percentage error in each property is presented in Fig. 6.17. The final correction fields constructed by the method are shown in Fig. 6.18.

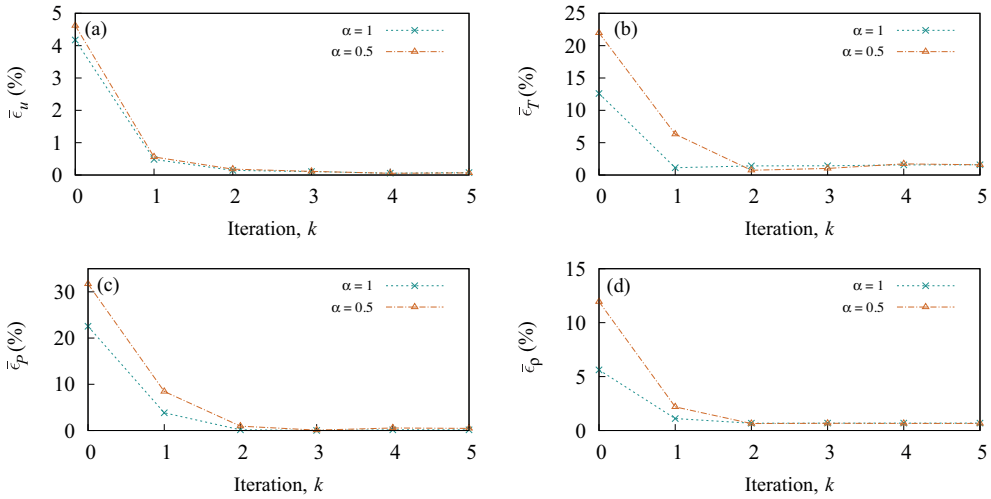

Figure 6.17: Hypersonic Couette flow: mean error in the hybrid a) velocity $\bar{\epsilon}_u$, b) temperature $\bar{\epsilon}_T$, c) pressure $\bar{\epsilon}_p$, and d) density $\bar{\epsilon}_\rho$ solutions for fully diffuse (α=1) and mixed specular-diffuse (α=0.5) walls; k=0 is the NSF solution.

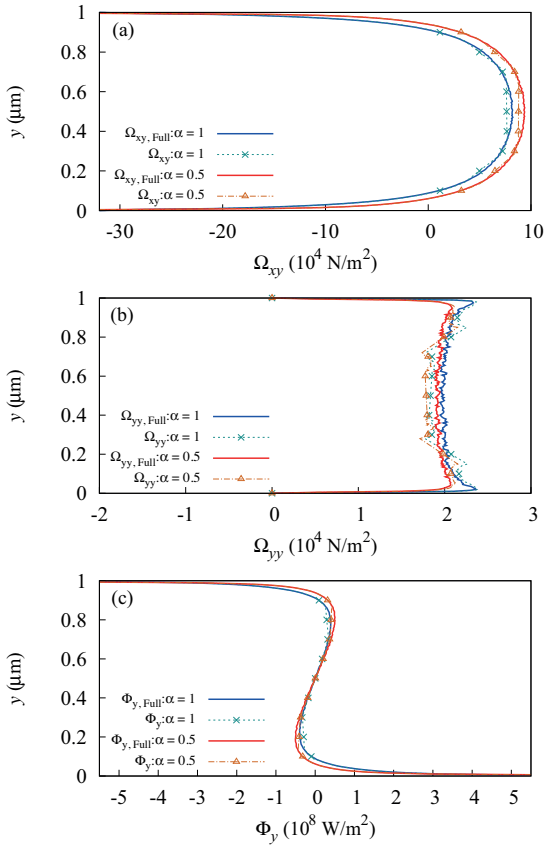

Figure 6.18: Hypersonic Couette flow: full DSMC and final hybrid (a) shear stress Ω_{xy}, (b) normal stress Ω_{yy}, and (c) heat flux Φ_y correction fields for fully diffuse (α=1) and mixed specular-diffuse (α=0.5) walls.

Figure 6.16 shows that, compared with fully diffuse walls, specular-diffuse walls result in greater velocity slip (and hence a larger momentum Knudsen layer), greater temperature jump, a higher maximum temperature in the centre of the flowfield, and a higher average pressure. The NSF solution (which is the same for both cases) is less accurate for the mixed specular-diffuse case than for the fully diffuse case. However, as the micro element configuration of the hybrid method is able to construct good approximations of all three correction fields, the final hybrid solution is equally accurate for both types of wall boundary; $\bar{\epsilon}_u$, $\bar{\epsilon}_p$, and $\bar{\epsilon}_\rho$ are less than 1%, and $\bar{\epsilon}_T$ is less than 2%.

6.4 Computational savings

As in the Fourier flow problem, the Couette flow test cases examined in this chapter perform the micro element and full DSMC simulations for the same number of time steps, i.e. $G_{\text{tot,Hyb}} = G_{\text{tot,Full}}$. Computational savings can then come only from spatial scale separation. The computational speed-up S offered by the hybrid method is calculated according to Eqs. (5.8) and (5.9). For each case, S is determined by the spatial extents of the near-wall micro elements, $H_{\text{lower}}=(H_{SR}+H_{RR})_{\text{lower}}$ and $H_{\text{upper}}=(H_{SR}+H_{RR})_{\text{upper}}$, and the spatial extents of any bulk micro elements, $H_{\text{bulk}}=(H_{SR}+2H_{RR})_{\text{bulk}}$.

Table 6.2 presents the speed-up from each of the micro element configurations considered for the supersonic test case of section 6.3.1. The initial configuration (with $\Pi = 2$ and $H_{\text{lower}} = H_{\text{upper}} = 10\lambda_l$) offers a modest speed-up of 1.72×. This value of S is, however, based on I=3 iterations. Although the temperature, pressure, and density fields are able to converge inside 3 iterations, the velocity field is unable to converge with this configuration, and the final results are quite poor.

Larger near-wall elements (i.e. $H_{\text{lower}} = H_{\text{upper}} = 20\lambda_l$) and additional bulk elements (with $H_{\text{bulk}} = 20\lambda_l$) enable the method to converge fully inside I=3 iterations, and improve the method's accuracy. Unfortunately, the increased computational cost that accompanies the increase in the micro resolution makes the hybrid method more expensive than the full DSMC simulations, i.e. $S < 1$ for $\Pi = 2$, 3, and 4; see Table 6.2. This is also true of the same three configurations applied to the hypersonic case of section 6.3.2 — even with convergence inside I=3 iterations, Table 6.3 shows that $S < 1$ for all three.

The spatial element extents presented in both Tables 6.2 and 6.3 are those at k=3 and subsequent iterations; note the changes from their original extents of 0.2 μm (not including the initial configuration applied to the supersonic case, which had original element extents of 0.1 μm). Comparing Tables 6.2 and 6.3, it is clear that the extents of the near-wall elements are smaller for the hypersonic case than for the supersonic case — this is because the variation of the density field ρ_{corr} is much more extreme in the hypersonic case, resulting in a higher density (and hence smaller local mean free path λ_l) close to the walls.

With the low spatial scale separation of these Couette flow test cases and the need to capture multiple constitutive correction fields, the hybrid method has been unable to provide computational savings for this flow problem (and still provide sufficient accuracy). However, the main purpose of these test cases was simply to validate the full coupling algorithm in 1D. As mentioned in the previous chapter, simulations of larger 2D/3D flow problems are expected to highlight any real computational advantages of the method.

Π	H_{lower} (μm)	H_{bulk} (μm)	H_{upper} (μm)	I	S
2 (initial)	0.095	-	0.095	3	1.72
2	0.19	-	0.19	3	0.87
3	0.19	0.205	0.19	3	0.57
4	0.19	0.205	0.19	3	0.43

Table 6.2: Computational speed-ups S for the four micro element configurations applied to the supersonic test case of section 6.3.1.

Π	H_{lower} (μm)	H_{bulk} (μm)	H_{upper} (μm)	I	S
2	0.17	-	0.17	3	0.99
3	0.17	0.285	0.17	3	0.55
4	0.17	0.27	0.17	3	0.41

Table 6.3: Computational speed-ups S for the three micro element configurations applied to the hypersonic test case of section 6.3.2.

6.5 Summary

The HMM-FWC continuum-DSMC method has been implemented for a micro Couette flow problem, where high wall speeds cause viscous heating in the gas. The method's full coupling algorithm (with one-way density coupling) has been demonstrated in 1D, considering simultaneous momentum and heat transport. The need to capture multiple constitutive corrections meant that the method's accuracy (compared with equivalent full DSMC simulations) depended heavily on the micro resolution, both close to the walls and in the bulk. Unfortunately, due to the low spatial scale separation of the test cases considered, the micro resolution required to obtain sufficient accuracy resulted in the hybrid method being more computationally expensive than the equivalent full DSMC simulations. Nevertheless, like the Fourier flow test cases of the previous chapter, these cases were chosen simply to validate the coupling strategy of the hybrid method, and greater computational speed-ups might be afforded when the method is applied to the larger 2D flow problem in the next chapter.

Chapter 7

Flow through a microscale crack: a 2D application

In this chapter, interest lies in modelling a rarefied gas flowing through a channel that has a high-aspect-ratio cross-section. Such channels can be encountered in MEMS devices, but are also representative of narrow cracks. Leakage from narrow cracks is fairly common in engineering systems like pressure vessels, pipes, and valves — as 'leak-before-break' analysis is commonly adopted as a form of safety assessment (and as an assessment of product reliability and performance), it is important that the leak rates through these cracks can be estimated accurately.

Consider the idealised crack-type channel geometry shown in Fig. 7.1 (where the cross-section of the flow is perfectly rectangular): the width W in the x-direction is much greater than the height H in the y-direction, i.e. $W/H \gg 1$; the streamwise length L is then much greater than the width.

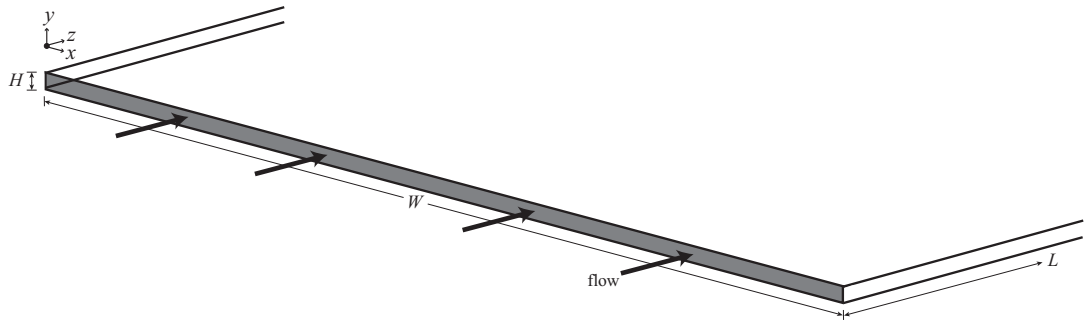

Figure 7.1: Schematic of an idealised 3D crack-type channel geometry.

To date, the study of gas flows in crack-type geometries has focused mainly on the continuum and slip regimes [124], with conventional fluid mechanics used to predict the flow behaviour. However, if the mean free path of the gas is of the same order as the crack opening (i.e. the height H), the flow will be non-equilibrium and predictions from conventional fluid mechanics will be poor. While the DSMC method could provide accurate flow predictions for such conditions, full DSMC simulations are likely to be prohibitively expensive for very high cross-sectional aspect-ratios. Existing DSMC studies [125, 126] have lessened this expense by assuming negligible variation of the flow in the direction of the width (essentially considering a 2D model in height and length), but this assumption will not hold true if there is significant geometric variation in the direction of the width. In reality, the cross-section of a crack will not be perfectly rectangular, for example, it could be more open at one end than the other; surface roughness or even surface defects could also have substantial effects on the flow behaviour.

A continuum-DSMC hybrid approach could be useful for modelling more realistic crack-type flows that see flow variation in the direction of the width. As the entire flowfield is near-wall, the popular DD framework would not be suitable. The point-wise coupling approach of the original HMM would also be inefficient as the micro resolution required over the crack/channel height would force the micro elements to overlap. While the IMM could exploit high scale separation in the streamwise z-direction, it would require each micro element to occupy the entire cross-sectional area and so could not exploit scale separation in the direction of the width. Our HMM-FWC method has the potential to deal effectively with the mixed degrees of spatial scale separation that can exist over the cross-section of such flows — it is able to exploit high scale separation in the direction of the width by distributing the micro elements across this dimension; simultaneously, it can cope with the low scale separation in the direction of the height by setting the micro elements to occupy the entire crack height. The higher the scale separation (i.e. the more gradual the variation of the flow properties) in the direction of the width, the more efficient our hybrid method will be.

To maintain simplicity and limit the expense of the full DSMC simulations that we use for validation, we assume that there is no flow variation in the streamwise z-

direction: periodicity is applied in this direction and an external acceleration a_z is used drive the flow. We then demonstrate the method's ability to cope with mixed spatial scale separation by applying it to the 2D cross-section of the channel only; see Fig. 7.2. Also for simplicity, we assume the channel cross-section to be perfectly rectangular, i.e. there is no geometric variation in the direction of the width. Although it is not necessarily representative of real crack flows, variation of the flow properties over the crack width is instead considered by imposing a temperature difference between the left and right walls of the channel, with an associated linear temperature gradient imposed on the lower and upper walls. A temperature difference between the lower and upper walls is also considered, with an associated linear temperature gradient imposed on the left and right walls — this could be representative of, for example, flow in a high-aspect-ratio microchannel heat exchanger. An isothermal test case and both of these non-isothermal test cases will demonstrate the hybrid coupling procedure (including simultaneous momentum and energy coupling) in 2D. Once again, the accuracy and the computational cost of the hybrid method will be compared with equivalent full DSMC simulations.

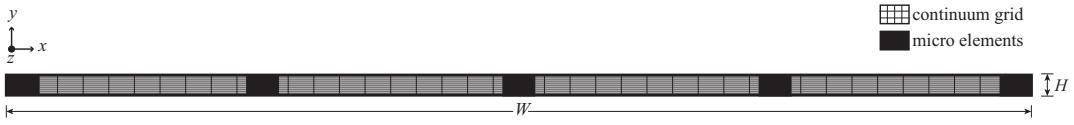

Figure 7.2: Schematic of the HMM-FWC framework applied to the high-aspect-ratio cross-section of a microchannel (flow is into the page); this acts as a representative geometry for a narrow crack. The continuum grid is applied over the entire cross-section and the DSMC particle simulations are performed in the micro elements.

More realistic geometric configurations could be considered in future work. The HMM-FWC method could also be applied in 3D, with the streamwise extents of the micro elements depending on the degree of scale separation in this direction. Alternatively, if the streamwise scale separation was very high, then significant computational savings could be achieved by combining the HMM-FWC method with the IMM: the HMM-FWC would comprise the 'micro elements' spaced intermittently along the streamwise direction of an IMM implementation.

7.1 Computational set-up

The computational set-up for this 2D flow problem is shown in Fig. 7.3: the 2D continuum grid applied over the entire cross-section is uniform, consisting of M_x macro nodes across the width and M_y macro nodes over the height (including nodes at the bounding walls); the horizontal and vertical node spacings are Δx and Δy, respectively. The micro elements span the entire height H. The element arrangement shown is an example, consisting of a 'side-wall' element at each of the left and right walls and one in the 'bulk' of the system. Each side-wall element comprises a single sampling region and a single relaxation region, while each bulk element consists of a single sampling region with a relaxation region on either side. All sampling regions are discretised into a grid of 2D measurement bins, while all relaxation regions are split into a grid of 2D control bins. Here, the measurement and control bins are set to have the same width δx and height δy.

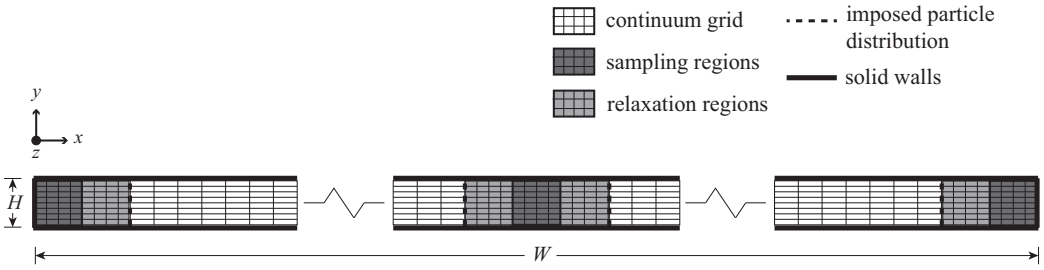

Figure 7.3: Schematic of the computational set-up for a 2D crack-type flow problem. The macro nodes are represented by the intersections of the horizontal and vertical lines of continuum grid. Note that the 2D bins of the sampling and relaxation regions are independent of the macro nodes.

7.2 Coupling algorithm

Compared with its implementation for 1D flow problems, the implementation of our hybrid method in 2D presents additional challenges: the continuum conservation equations become more complex; the micro elements must be constrained by applying boundary conditions within 2D relaxation regions; and the interpolations of the constitutive correction fields between the sampling regions must be 2D.

Once again, constant reference values of the transport properties are adopted, μ_{ref}

and κ_{ref}, with the constitutive corrections capturing the true variation of these properties as the iterations proceed. With subsonic flow and periodicity in the streamwise z-direction, mass conservation is satisfied automatically, and the transport of momentum in the x- and y-directions is assumed negligible, i.e. the gas pressure and density are assumed constant over the flow cross-section. In the streamwise z-direction, the modified momentum equation (4.6) is reduced to,

$$\mu_{\text{ref}} \left(\frac{\partial^2 w}{\partial x^2} + \frac{\partial^2 w}{\partial y^2} \right) + \mathrm{f}_z + \Psi = 0, \tag{7.1}$$

where w is the streamwise velocity component, f_z is the streamwise external force per unit volume, and the 'overall momentum correction' term Ψ is the divergence of the constitutive shear stress correction, i.e.

$$\Psi = \frac{\partial \Omega_{zx}}{\partial x} + \frac{\partial \Omega_{zy}}{\partial y}. \tag{7.2}$$

In isothermal conditions, this momentum equation will be sufficient to describe the macroscopic flow variation. The two constitutive stress corrections are computed from the sampling regions of the micro elements by rearranging Eq. (4.4), i.e.

$$\Omega_{zx} = \tau_{zx} - \mu_{\text{ref}} \frac{dw}{dx}, \tag{7.3}$$

and

$$\Omega_{zy} = \tau_{zy} - \mu_{\text{ref}} \frac{dw}{dy}. \tag{7.4}$$

In non-isothermal conditions (i.e. if there is a temperature difference between the bounding walls of the cross-section), the transport of energy must also be considered. Using the product rule, the modified energy equation (4.7) is reduced to,

$$\begin{aligned} \mu_{\text{ref}} \left[w \left(\frac{\partial^2 w}{\partial x^2} + \frac{\partial^2 w}{\partial y^2} \right) + \left(\frac{\partial w}{\partial x} \right)^2 + \left(\frac{\partial w}{\partial y} \right)^2 \right] \\ + \kappa_{\text{ref}} \left(\frac{\partial^2 T}{\partial x^2} + \frac{\partial^2 T}{\partial y^2} \right) + \Upsilon = 0, \end{aligned} \tag{7.5}$$

where the 'overall energy correction term' Υ encompasses both the constitutive stress

and heat flux corrections,

$$\Upsilon = \Omega_{zx}\frac{\partial w}{\partial x} + \Omega_{zy}\frac{\partial w}{\partial y} + w\left(\frac{\partial \Omega_{zx}}{\partial x} + \frac{\partial \Omega_{zy}}{\partial y}\right) - \left(\frac{\partial \Phi_x}{\partial x} + \frac{\partial \Phi_y}{\partial y}\right). \tag{7.6}$$

The two constitutive heat flux corrections are computed from the sampling regions by rearranging Eq. (4.5), i.e.

$$\Phi_x = q_x + \kappa_{\text{ref}}\frac{\partial T}{\partial x}, \tag{7.7}$$

and

$$\Phi_y = q_y + \kappa_{\text{ref}}\frac{\partial T}{\partial y}. \tag{7.8}$$

Therefore, for the non-isothermal cases of this flow problem, the hybrid coupling algorithm is as follows[1].

(0) With no corrections ($\Psi = \Upsilon = 0$), five-point finite difference approximations are used to represent Eqs. (7.1) and (7.5), i.e.

$$\mu_{\text{ref}}\left[\frac{(w_{i-1,j} - 2w_{i,j} + w_{i+1,j})}{\Delta x^2} + \frac{(w_{i,j-1} - 2w_{i,j} + w_{i.j+1})}{\Delta y^2}\right] + \mathrm{f}_{z_{i,j}} = 0, \tag{7.9}$$

and,

$$\begin{aligned}
&\mu_{\text{ref}} w_{i,j}\left[\frac{(w_{i-1,j} - 2w_{i,j} + w_{i+1,j})}{\Delta x^2} + \frac{(w_{i,j-1} - 2w_{i,j} + w_{i,j+1})}{\Delta y^2}\right] \\
+\mu_{\text{ref}}&\left[\frac{(w_{i+1,j}^2 - 2w_{i-1,j}w_{i+1,j} + w_{i-1,j}^2)}{4\Delta x^2} + \frac{(w_{i,j+1}^2 - 2w_{i,j-1}w_{i,j+1} + w_{i,j-1}^2)}{4\Delta y^2}\right] \\
+\kappa_{\text{ref}}&\left[\frac{(T_{i-1,j} - 2T_{i,j} + T_{i+1,j})}{\Delta x^2} + \frac{(T_{i,j-1} - 2T_{i,j} + T_{i,j+1})}{\Delta y^2}\right] = 0,
\end{aligned} \tag{7.10}$$

respectively, where i=2, 3, ...(M_x −1) and j=2, 3, ...(M_y −1). With no velocity slip at the solid walls, solving Eq. (7.9) obtains the initial NSF streamwise velocity field w_{NSF}. This NSF velocity field is then used in Eq. (7.10); with no temperature jump, solution of this equation produces the NSF temperature field T_{NSF}.

(1) All micro elements are initialised at equilibrium; the velocity and temperature

[1]The same algorithm is implemented for isothermal cases by neglecting the conservation of energy (and hence the calculation of the heat flux corrections) and the imposition of the continuum temperature field over the relaxation regions.

for initialisation are obtained by averaging over a sub-region of the continuum solution that corresponds to the element.

Each element is then constrained as follows:

(a) The local continuum velocity and temperature fields are imposed across each relaxation region by implementing the velocity and temperature state controllers in each 2D control bin.

(b) A Maxwellian particle distribution is imposed at the outer boundaries of each relaxation region via a diffuse reflection boundary at the local continuum velocity and temperature.

(c) An external acceleration $a_z = \mathrm{f}_z/\rho$ is applied to all particles in the element (in both the sampling and relaxation regions).

(2) The DSMC algorithm is executed in each micro element for a transient start-up period, followed by a subsequent steady-state averaging period.

(3) The time-averaged macroscopic flow properties are extracted from the 2D measurement bins of all sampling regions. The gas properties at the walls are also extracted from the wall-adjacent bin faces of all sampling regions.

(4) The shear stress corrections in each measurement bin and at the walls are computed according to Eqs. (7.3) and (7.4); the velocity gradients are approximated using finite difference representations, e.g. for a bin $b_{i,j}$,

$$\Omega_{zx_{b_{i,j}}} = \tau_{zx_{b_{i,j}}} - \mu_{\mathrm{ref}} \frac{(w_{b_{i+1,j}} - w_{b_{i-1,j}})}{2\delta x}, \tag{7.11}$$

and

$$\Omega_{zy_{b_{i,j}}} = \tau_{zy_{b_{i,j}}} - \mu_{\mathrm{ref}} \frac{(w_{b_{i,j+1}} - w_{b_{i,j-1}})}{2\delta y}. \tag{7.12}$$

Similarly, the heat flux corrections in each measurement bin and at each wall are computed from Eqs. (7.7) and (7.8); again, finite difference approximations are used for the temperature gradients, e.g.

$$\Phi_{x_{b_{i,j}}} = q_{x_{b_{i,j}}} + \kappa_{\mathrm{ref}} \frac{(T_{b_{i+1,j}} - T_{b_{i-1,j}})}{2\delta x}, \tag{7.13}$$

and

$$\Phi_{y_{b_{i,j}}} = q_{y_{b_{i,j}}} + \kappa_{\text{ref}} \frac{(T_{b_{i,j+1}} - T_{b_{i,j-1}})}{2\delta y}. \tag{7.14}$$

(5) Global approximations of Ω_{zx}, Ω_{zy}, Φ_x, and Φ_y over the entire cross-section are obtained by interpolating between the local sampling region corrections.

To approximate the global boundary information at all wall surfaces, the local gas boundary velocities and temperatures extracted in Step (3) are interpolated linearly across the channel width.

(6) Now with the global constitutive corrections, Eqs. (7.1) and (7.5) can again be approximated by a five-point finite difference scheme, i.e.

$$\mu_{\text{ref}} \left[\frac{(w_{i-1,j} - 2w_{i,j} + w_{i+1,j})}{\Delta x^2} + \frac{(w_{i,j-1} - 2w_{i,j} + w_{i,j+1})}{\Delta y^2} \right] + \mathrm{f}_{z_{i,j}} + \Psi_{i,j} = 0, \tag{7.15}$$

where the overall momentum correction is approximated as,

$$\Psi_{i,j} = \frac{(\Omega_{zx_{i+1,j}} - \Omega_{zx_{i-1,j}})}{2\Delta x} + \frac{(\Omega_{zy_{i,j+1}} - \Omega_{zy_{i,j-1}})}{2\Delta y}, \tag{7.16}$$

and

$$\begin{aligned} &\mu_{\text{ref}} w_{i,j} \left[\frac{(w_{i-1,j} - 2w_{i,j} + w_{i+1,j})}{\Delta x^2} + \frac{(w_{i,j-1} - 2w_{i,j} + w_{i,j+1})}{\Delta y^2} \right] \\ &+\mu_{\text{ref}} \left[\frac{(w_{i+1,j}^2 - 2w_{i-1,j}w_{i+1,j} + w_{i-1,j}^2)}{4\Delta x^2} + \frac{(w_{i,j+1}^2 - 2w_{i,j-1}w_{i,j+1} + w_{i,j-1}^2)}{4\Delta y^2} \right] \\ &+\kappa_{\text{ref}} \left[\frac{(T_{i-1,j} - 2T_{i,j} + T_{i+1,j})}{\Delta x^2} + \frac{(T_{i,j-1} - 2T_{i,j} + T_{i,j+1})}{\Delta y^2} \right] + \Upsilon_{i,j} = 0, \end{aligned} \tag{7.17}$$

where the overall energy correction is,

$$\begin{aligned} \Upsilon_{i,j} = {} & \Omega_{zx_{i,j}} \frac{(w_{i+1,j} - w_{i-1,j})}{2\Delta x} + \Omega_{zy_{i,j}} \frac{(w_{i,j+1} - w_{i,j-1})}{2\Delta y} \\ & + w_{i,j} \left[\frac{\left(\Omega_{zx_{i+1,j}} - \Omega_{zx_{i-1,j}}\right)}{2\Delta x} + \frac{\left(\Omega_{zy_{i,j+1}} - \Omega_{zy_{i,j-1}}\right)}{2\Delta y} \right] \\ & - \left[\frac{\left(\Phi_{x_{i+1,j}} - \Phi_{x_{i-1,j}}\right)}{2\Delta x} + \frac{\left(\Phi_{y_{i,j+1}} - \Phi_{y_{i,j-1}}\right)}{2\Delta y} \right], \end{aligned} \tag{7.18}$$

where i=2, 3, ...$(M_x - 1)$ and j=2, 3, ...$(M_y - 1)$. Using the updated global boundary velocity information, solution of Eq. (7.15) provides the new, corrected streamwise velocity field w_{corr} over the cross-section. Subsequently, using this corrected velocity field and the updated global boundary temperature information to solve Eq. (7.17) produces the new, corrected temperature field T_{corr}.

(7) Using w_{corr} and T_{corr} to replace w_{NSF} and T_{NSF}, the algorithm is repeated from Step (1) until both fields converge to within user-defined tolerances.

7.3 Results

All of the test cases considered in this chapter are based on the same channel geometry: with a width W of 200 μm and a height H of 0.1 μm, the aspect-ratio of the cross-section is 2000. Greater macroscopic resolution is needed in the direction of the height than in the direction of the width; $M_y = 26$ macro nodes over the height (including nodes at the lower and upper walls) give a vertical node spacing $\Delta y = 4$ nm, while $M_x = 801$ nodes across the width (including nodes at the left and right walls) give a horizontal spacing $\Delta x = 250$ nm. The micro elements of the hybrid method capture a somewhat higher resolution — the height of each measurement/control bin δy is equal to the vertical macro node spacing of 4 nm, but the bin width δx is 50 nm.

Argon is the working gas in all of the cases of this chapter (with the VHS parameters stated in section 5.3). As mentioned above, both isothermal and non-isothermal cases are considered: in the isothermal case, the gas and the fully diffuse bounding walls are maintained at T_{av}= T_{wall}= 273 K; in the non-isothermal cases, temperature differences are present between the diffuse bounding walls but the average gas temperature remains at T_{av}= 273 K. Therefore, based on this temperature, reference values $\mu_{\text{ref}} = 2.1\times10^{-5}$ kg/ms and $\kappa_{\text{ref}} = 0.0164$ W/mK [121] are adopted for all cases.

The 1D test cases of the previous chapters were restricted (due to the low spatial scale separation in the problems investigated) to low values of Kn_{gl}. All of the cases in this chapter are, however, based on Kn_{gl}= 1. The global characteristic dimension for this crack-type geometry is the height H, i.e. Kn_{gl}= λ_{gl}/H, and so, from Eq. (2.9), an average gas density ρ_{av} of 0.85 kg/m^3 is set to obtain λ_{gl}= 0.1 μm.

To generate low-speed flow in the z-direction, all of the micro DSMC simulations of the hybrid method (and the comparable full DSMC simulations) apply an acceleration a_z of $10{\times}10^{10}$ m/s^2 to each particle; this corresponds to a body forcing $\mathrm{f}_z = \rho_{\mathrm{av}} a_z = 8.5{\times}10^{10}$ N/m^3. It is common practice to apply such large and unrealistic forcing in particle/molecular simulations, so that the resulting signal-to-noise ratio is manageable. For both isothermal and non-isothermal conditions, the conventional no-slip NSF solution of this flow problem predicts a maximum flow velocity of approximately 5 m/s in the centre of the cross-section, i.e. $Ma \approx 0.016$.

All of the micro and full DSMC simulations performed in this chapter use a micro time step $\delta t = 5{\times}10^{-12}$ s. Steady-state conditions are reached inside a transient start-up period of 0.6 million time steps, and a further 1.4 million time steps are then performed to reduce the statistical scatter in the measured flow properties.

For this 2D problem, the convergence of each property field s produced from the hybrid method (i.e. for the isothermal case, $s{=}w$; for the non-isothermal cases, $s{=}w, T$) is monitored at each iteration k according to,

$$\zeta_s{}^k = \frac{1}{M_x M_y} \sum_{j}^{M_y} \sum_{i}^{M_x} \left| \frac{s^k_{\mathrm{corr}_{i,j}} - s^{k-1}_{\mathrm{corr}_{i,j}}}{s^k_{\mathrm{corr}_{i,j}}} \right| \leq \zeta_{\mathrm{tol}_s}, \tag{7.19}$$

where $i{=}1, 2, ...M_x$ and $j{=}1, 2, ...M_y$. In this chapter, tolerance values of $\zeta_{\mathrm{tol}_w}{=}3{\times}10^{-3}$ and $\zeta_{\mathrm{tol}_T}{=}3{\times}10^{-4}$ are adopted. The accuracy of each hybrid property field is then measured by the mean percentage error compared with the full DSMC solution, i.e.

$$\bar{\epsilon}^k_s = \frac{1}{M_x M_y} \sum_{j}^{M_y} \sum_{i}^{M_x} \left[\frac{s_{\mathrm{Full}_{i,j}} - s^k_{\mathrm{corr}_{i,j}}}{s_{\mathrm{Full,\,max}} - s_{\mathrm{Full,\,min}}} \times 100\% \right]. \tag{7.20}$$

As both property fields are expected to vary considerably over the cross-section, the denominator of this expression is the range of the full DSMC solution. These errors will depend on the construction of Ω_{zx}, Ω_{zy}, Φ_x, and Φ_y, which should ideally match the full DSMC corrections $\Omega_{zx,\,\mathrm{Full}}$, $\Omega_{zy,\,\mathrm{Full}}$, $\Phi_{x,\,\mathrm{Full}}$, and $\Phi_{y,\,\mathrm{Full}}$ that are calculated by using the full DSMC property fields in Eqs. (7.3), (7.4), (7.7), and (7.8), respectively.

When investigating fluid flows through cracks and general microchannels, interest often lies in predicting the mass flow rate $\dot{m}$ along the crack/channel length. For the

2D test cases of this chapter, the mass flow rates from each solution (i.e. the NSF approach, the hybrid method, and the full DSMC simulation) are approximated using a surface integral,

$$\dot{m} = \int_A \rho w dA \approx \sum_j^{M_y} \sum_i^{M_x} \left[\rho_{i,j} w_{i,j} \Delta x \Delta y\right], \tag{7.21}$$

where i=1, 2, ...M_x and j=1, 2, ...M_y. As well as using the different velocity fields from each solution (w_{NSF}, w_{Hyb}= w_{corr}, and w_{Full}), this mass flow rate calculation is also based on the different density fields: ρ_{NSF} and ρ_{Hyb} are simply equal to the average density ρ_{av} everywhere, while ρ_{Full} is the density field predicted from the full DSMC simulation. For these low-speed test cases, ρ_{Full} will vary only slightly from the value of ρ_{av}.

7.3.1 Isothermal flow

The no-slip NSF description of this isothermal flow problem (where T_{wall}=T_{av}=273 K) predicts a velocity field w_{NSF} with a maximum of approximately 5 m/s (i.e. $Ma \approx 0.016$) in the centre of the cross-section; this is shown in Fig. 7.4(a). From Eq. (7.3), this velocity field gives a mass flow rate prediction $\dot{m}_{\text{NSF}}$= 5.74×10^{-11} kg/s. A full DSMC simulation of this same problem predicts a streamwise velocity w_{Full} with a much larger magnitude, and a maximum of approximately 50 m/s (i.e. $Ma \approx 0.16$) in the centre of the cross-section, as shown in Fig. 7.4(b). The mass flow rate prediction based on w_{Full} and ρ_{Full} is also significantly larger than the NSF prediction, with a value of $\dot{m}_{\text{Full}}$= 73.6×10^{-11} kg/s.

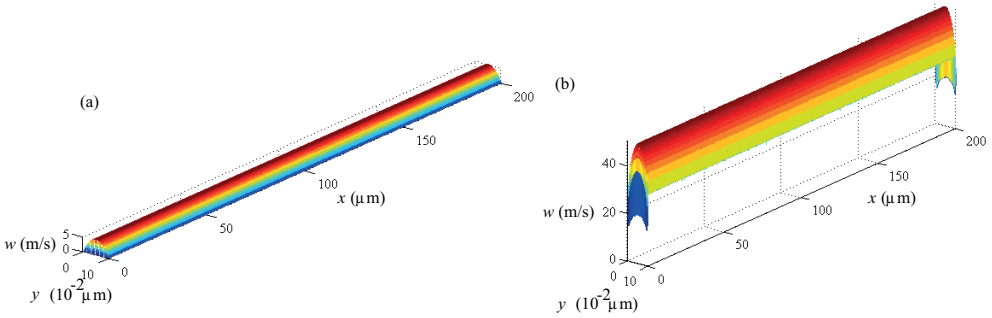

Figure 7.4: Isothermal crack flow: surface plot of (a) the NSF velocity field w_{NSF} and (b) the full DSMC velocity field w_{Full}.

The aim of the hybrid method is to iterate from the NSF solution towards the 'correct' full DSMC solution at a computational cost that is less than that of the full DSMC simulation. When implementing the method for this case, two micro element configurations are considered: the first uses only two side-wall elements, at the left and right walls (i.e. $\Pi = 2$); the second considers the addition of a bulk element (i.e. $\Pi = 3$) that is centred at x=0.5W. Following some initial testing, horizontal extents of W_{SR}=W_{RR}=5λ_l were deemed sufficient for the sampling and relaxation regions, with both adapting dynamically at each iteration depending on λ_l. The interpolations of the local shear stress correction fields and boundary velocities across the width are linear.

For both element configurations, the velocity fields produced from the hybrid method were able to converge within 6 iterations; see Fig. 7.5. To compare the final hybrid velocity fields with the NSF and full DSMC velocity fields shown in Fig. 7.4, the velocity profiles across the channel width at the centreline of its height (i.e. at y=0.5H) are depicted in Fig. 7.6(a). Similarly, the velocity profiles over the height at the centreline of the width (i.e. at x=0.5W) are shown in Fig. 7.6(b). The solution from both hybrid configurations is shown to agree almost exactly with the the full DSMC solution.

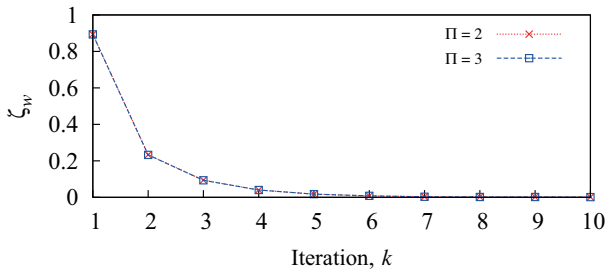

Figure 7.5: Isothermal crack flow: convergence of w_{corr} for both configurations.

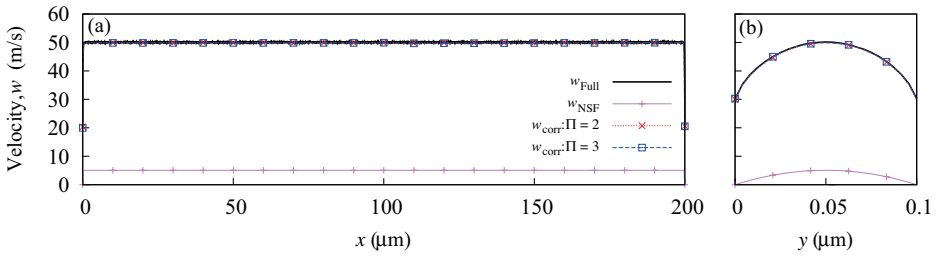

Figure 7.6: Isothermal crack flow: the NSF, full DSMC and final hybrid velocity solutions, (a) across the width at $y = 0.5H$, and (b) over the height at $x = 0.5W$, for both configurations.

The mean percentage error $\bar{\epsilon}_w$ presented in Fig. 7.7 confirms that the hybrid method is able to achieve a major increase in accuracy over the NSF solution — while the NSF velocity solution has a mean error of 112.98%, the hybrid method is able to reduce this to below 1% after 6 iterations, for both $\Pi = 2$ and $\Pi = 3$. The same is true of the mass flow rate: compared with the full DSMC prediction, the NSF prediction has an error of 92.2%; contrastingly, both hybrid configurations are able to predict $\dot{m}_{\text{Hyb}}$= 73.7×10^{-11} kg/s, equating to an error of only 0.14%.

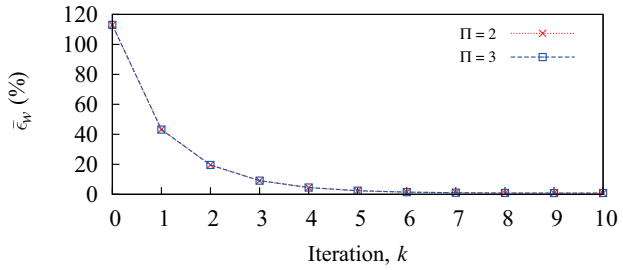

Figure 7.7: Isothermal crack flow: mean error $\bar{\epsilon}_w$ in the hybrid velocity field for both configurations; k=0 represents the initial NSF solution.

For this flow problem, there is no increase in accuracy with the addition of a bulk micro element. This is because both configurations are able to construct good approximations of the constitutive stress correction fields; the corrections computed from the full DSMC solution are depicted in Fig. 7.8. The excellent agreement between the final corrections approximated by the hybrid method and these full DSMC corrections is shown by the profiles across the width at y=0.5H and over the height at x=0.5W in Fig. 7.9. Although two side-wall elements are sufficient for this case, this may not be true if there is variation in the flow properties across the width.

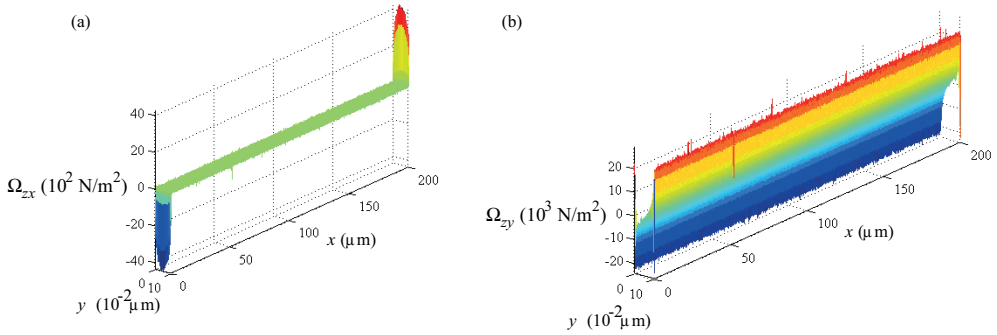

Figure 7.8: Isothermal crack flow: surface plots of the full DSMC stress correction fields, (a) Ω_{zx}, and (b) Ω_{zy}.

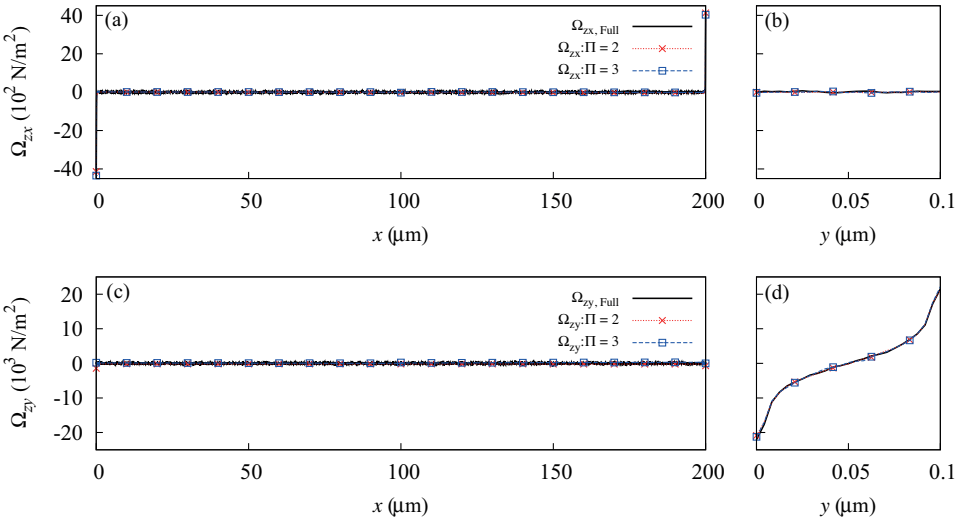

Figure 7.9: Isothermal crack flow: full DSMC and final hybrid stress corrections, Ω_{zx} (a) across the width at y=0.5H and (b) over the height at x=0.5W, and Ω_{zy} (c) across the width at y=0.5H and (b) over the height at x=0.5W, for both configurations.

The computational savings afforded by the hybrid method for this test case and the non-isothermal test cases in the following sub-sections will be discussed in section 7.4.

7.3.2 Non-isothermal flow: a temperature difference across the width

A temperature difference between the left and right walls is now considered: the left wall is set at T_{left}= 223 K while the right wall is set at T_{right}= 323 K; the associated linear temperature gradient is then imposed on the lower and upper walls.

The NSF velocity and temperature solutions for this case are presented in Fig. 7.10(a) and (c), respectively; w_{NSF} does not incorporate the effects of the temperature difference on the flow velocity, while T_{NSF} assumes a linear temperature variation in the x-direction. Once again, the streamwise velocity w_{Full} predicted from a full DSMC simulation has a much larger magnitude, as shown in Fig. 7.10(b); this velocity field is somewhat skewed, with a larger velocity at the colder left wall and a maximum of approximately 55 m/s (i.e. $Ma \approx 0.18$). Although it is difficult to observe from Fig. 7.10(d), the full DSMC temperature field T_{Full} features small thermal Knudsen layers at the left and right walls.

With variation of the flow properties in the x-direction, the hybrid method is performed with two side-wall elements and one bulk element at x=0.5W (i.e. $\Pi = 3$).

Each sampling and relaxation region has a horizontal extent W_{SR}=W_{RR}=5λ_l, and the interpolations of the local correction fields and boundary properties are linear.

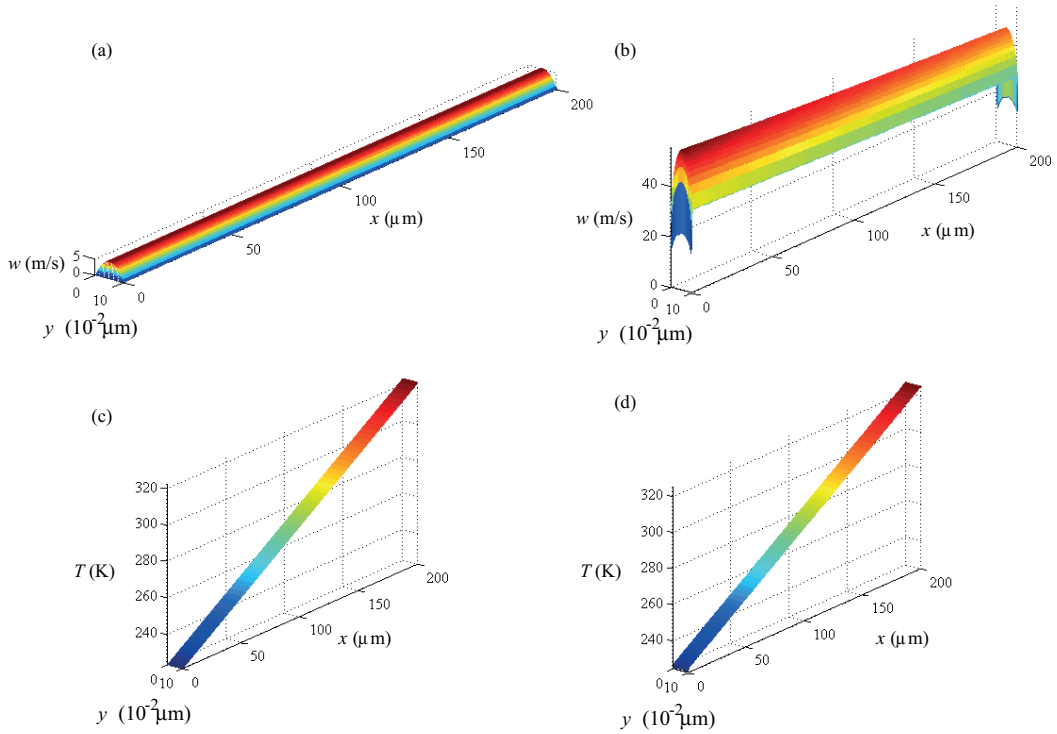

Figure 7.10: Crack flow with a temperature difference across the width: surface plot of (a) the NSF velocity field w_{NSF}, (b) the full DSMC velocity field w_{Full}, (c) the NSF temperature field T_{NSF}, and (d) the full DSMC temperature field T_{Full}.

Using this element configuration, Fig. 7.11 shows that both the velocity and temperature fields produced from the method reach convergence inside 6 iterations. The final hybrid velocity and temperature fields are then compared with the full DSMC fields by observing the profiles across the channel width at y=0.5H, and over the height at x=0.5W; see Fig. 7.12. For both properties, the method predicts fields that match almost exactly those obtained from the full DSMC simulation, capturing the skew of the velocity in the x-direction, and the thermal Knudsen layers at the side walls.

As for the isothermal case, the hybrid method provides a substantial reduction in the mean error $\bar{\epsilon}_w$: the initial NSF solution has a mean error value of 97.59%; the hybrid method reduces this to only 0.81% inside 5 iterations, as shown in Fig. 7.13(a). Although the initial NSF temperature solution has a mean error of only 3.11%, the method is able to reduce this to 0.25% inside 6 iterations; see Fig. 7.13(b).

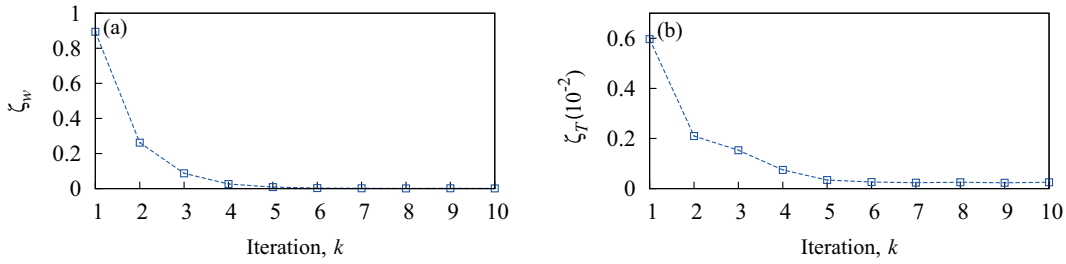

Figure 7.11: Crack flow with a temperature difference across the width: convergence of (a) w_{corr} and (b) T_{corr}.

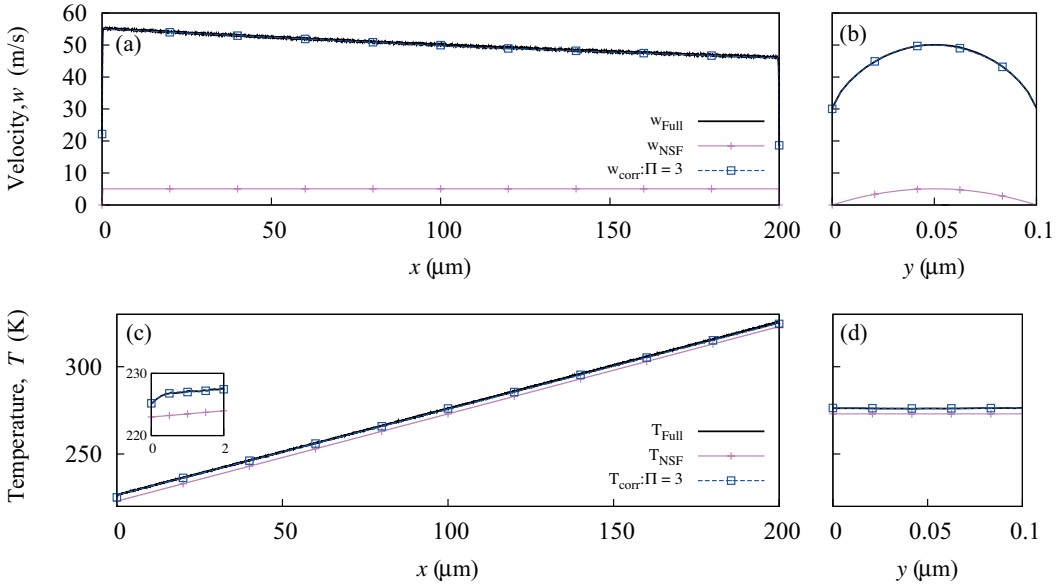

Figure 7.12: Crack flow with a temperature difference across the width: the NSF, full DSMC and final hybrid velocity solutions (a) across the width at y=0.5H and (b) over the height at x=0.5W, and temperature solutions (c) across the width at y=0.5H and (d) over the height at x=0.5W.

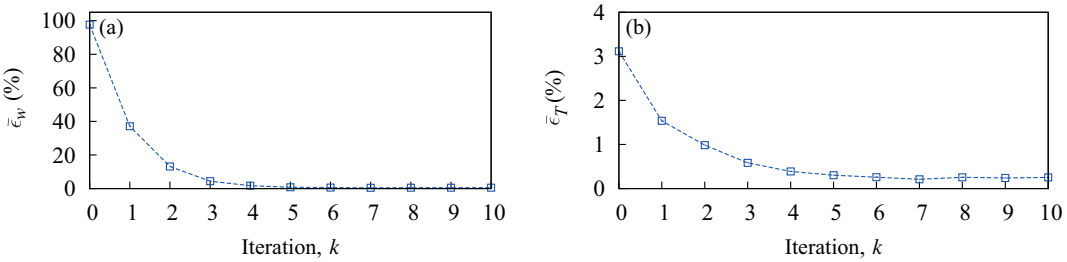

Figure 7.13: Crack flow with a temperature difference across the width: mean error in the hybrid (a) velocity and (b) temperature fields at each iteration; k=0 represents the NSF solution.

As well as providing highly accurate predictions of the velocity and temperature fields over the cross-section, the hybrid method is also able to provide an excellent approximation of the mass flow rate: compared with the full DSMC prediction for this test case, $\dot{m}_{\text{Full}}$= 73.97×10^{-11} kg/s, the hybrid prediction of $\dot{m}_{\text{Hyb}}$= 74.1×10^{-11} kg/s has an error of only 0.18%. On the other hand, the prediction from the NSF description (which is the same as that for the isothermal case, $\dot{m}_{\text{NSF}}$=5.74×10^{-11} kg/s) equates to an error of 92.24%.

Such high accuracy is achieved as this hybrid configuration is able to generate good approximations of all four constitutive correction fields; the full DSMC correction fields are presented in Fig. 7.14. The final correction fields constructed by the hybrid method are compared against these full DSMC corrections in Fig. 7.15, which shows a comparison of the profiles across the width at y=0.5H and over the height at x=0.5W for all four corrections. Note that substantial scatter makes it difficult to observe the true variation of the heat flux correction fields. Scatter in the bulk micro element also results in a discontinuity in the hybrid method's approximation of Φ_x near the centre of the channel width, which can be seen in Fig. 7.15(c). Fortunately, the effects of this discontinuity are not significant.

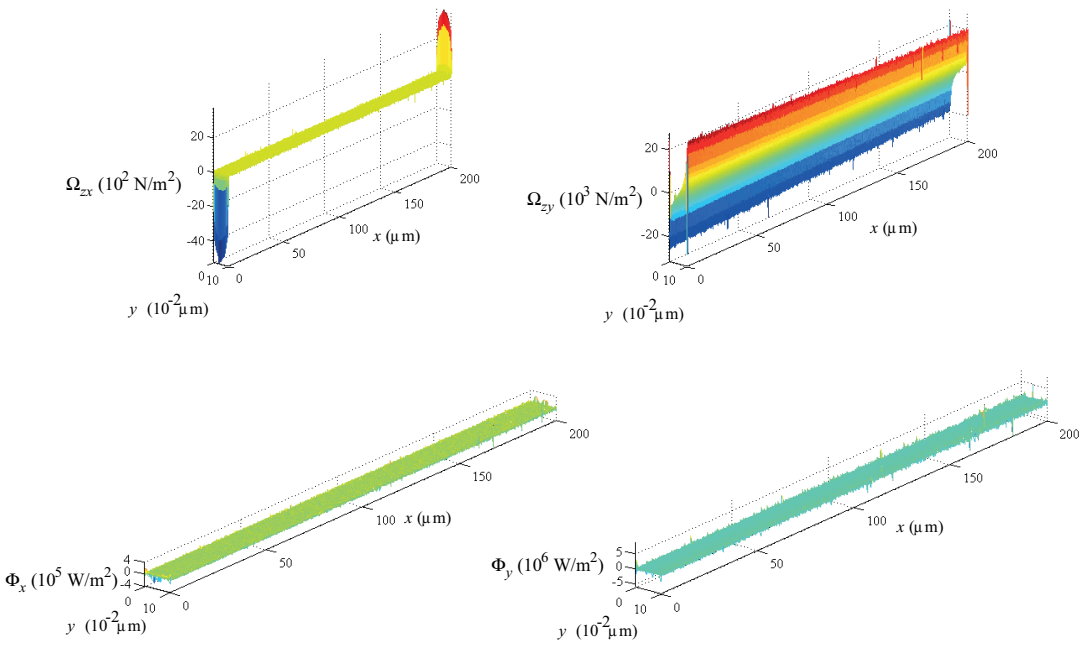

Figure 7.14: Crack flow with a temperature difference across the width: surface plot of the full DSMC (a) Ω_{zx}, (b) Ω_{zy}, (c) Φ_x, and (b) Φ_y correction fields.

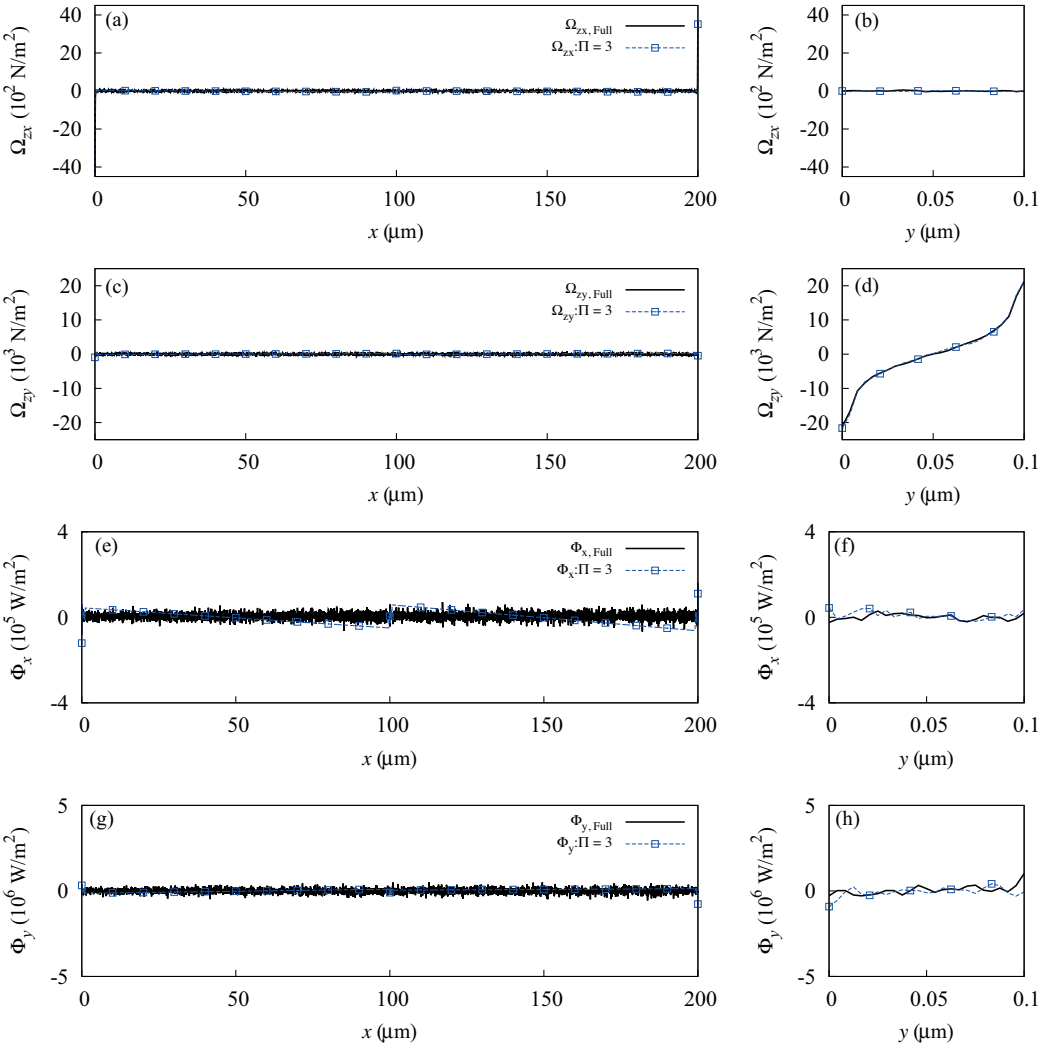

Figure 7.15: Crack flow with a temperature difference across the width: full DSMC and final shear stress corrections, Ω_{zx} over (a) y=0.5H and (b) x=0.5W, and Ω_{zy} over (c) y=0.5H and (d) x=0.5W. Similarly, the full DSMC and final heat flux corrections, Φ_x over (e) y=0.5H and (f) x=0.5W, and Φ_y over (g) y=0.5H and (h) x=0.5W.

7.3.3 Non-isothermal flow: a temperature difference over the height

To test the method further, the same temperature difference is now imposed between the lower and upper walls, i.e. the lower wall is set at T_{lower}= 223 K while the upper wall is set at T_{upper}= 323 K. A linear temperature gradient between these two temperatures is imposed on the left and right walls.

As in the non-isothermal test case of the previous sub-section, the NSF prediction of the streamwise velocity w_{NSF} for this problem does not incorporate the effects of the

temperature difference on the flow velocity, and the temperature field T_{NSF} assumes a linear temperature variation, now in the y-direction; see Figs. 7.16(a) and (c), respectively. Again, the velocity field w_{Full} predicted from a full DSMC simulation has a much larger magnitude; although it is difficult to see in Fig. 7.16(b), this field is somewhat skewed in the y-direction, with a maximum velocity of approximately 51 m/s (i.e. $Ma \approx 0.17$) closer to the hotter upper wall. The full DSMC temperature field T_{Full} depicted in Fig. 7.16(d) shows significant temperature jump and substantial thermal Knudsen layers at the lower and upper walls.

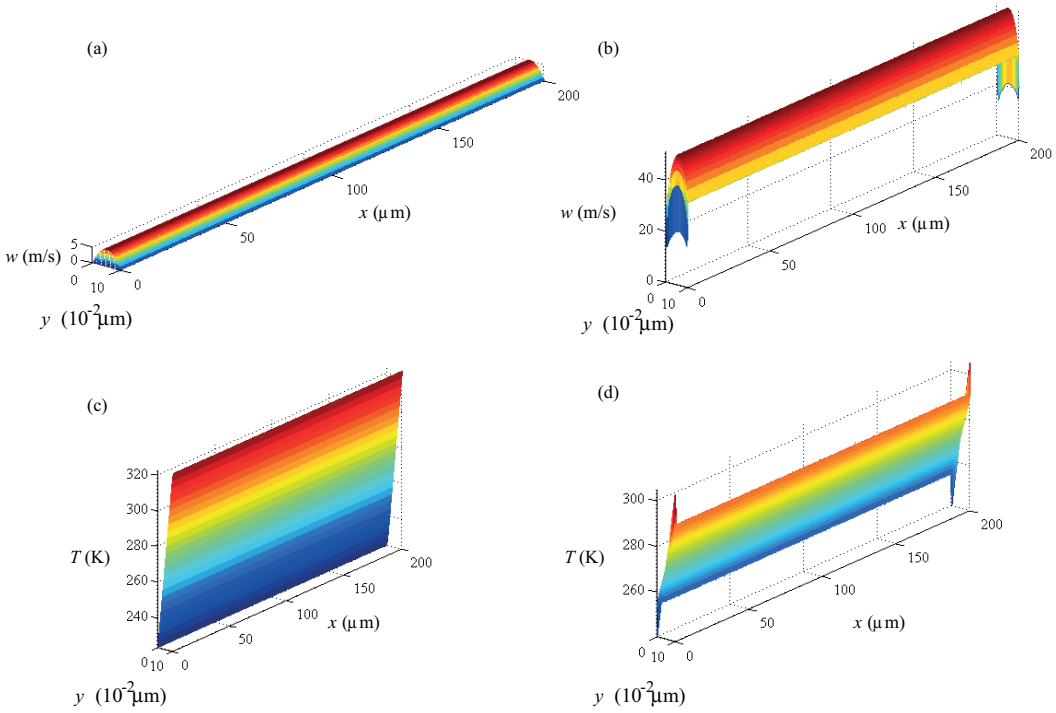

Figure 7.16: Crack flow with a temperature difference over the height: surface plot of (a) the NSF velocity field w_{NSF}, (b) the full DSMC velocity field w_{Full}, (c) the NSF temperature field T_{NSF}, and (d) the full DSMC temperature field T_{Full}.

The hybrid micro element configuration that was used to simulate the previous non-isothermal case is used again (i.e. Π=3 and W_{SR}=W_{RR}=5λ_l), and linear interpolations provide approximations of the global correction fields and boundary properties. With this configuration, the hybrid velocity and temperature solutions converge inside 7 iterations, as shown in Fig. 7.17. The velocity and temperature profiles depicted in Fig. 7.18 confirm that the method is, once again, able to predict property fields that agree closely with those of the full DSMC simulation: the hybrid velocity field is able

to model the skew of the velocity in the y-direction, and the hybrid temperature field is able to capture the large thermal Knudsen layers at the lower and upper walls.

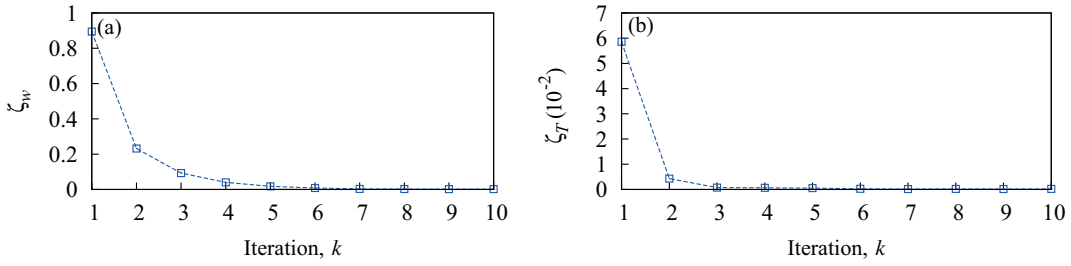

Figure 7.17: Crack flow with a temperature difference over the height: convergence of (a) w_{corr} and (b) T_{corr}.

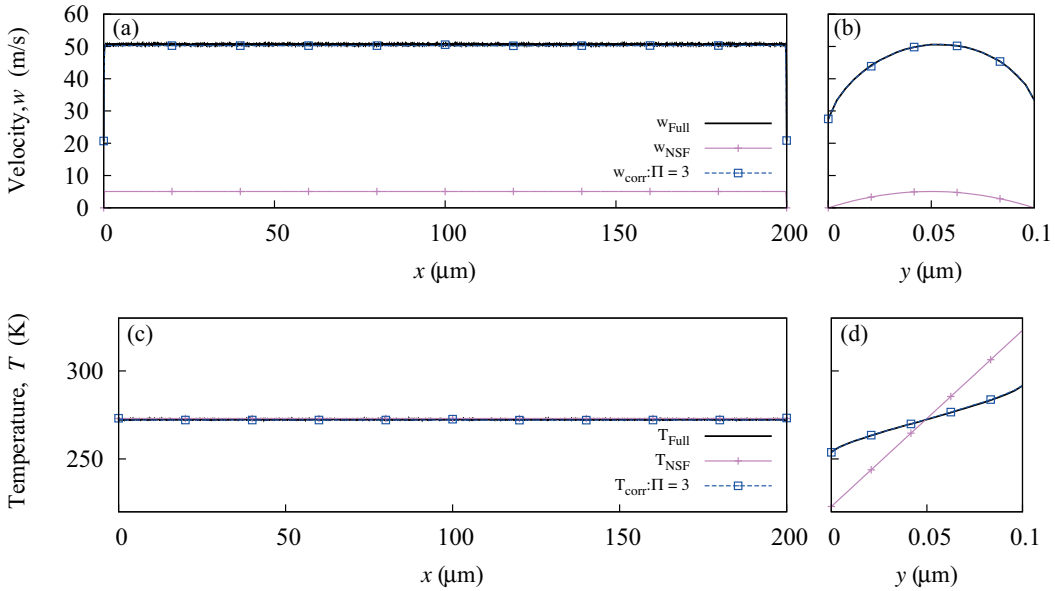

Figure 7.18: Crack flow with a temperature difference over the height: the NSF, full DSMC and final hybrid velocity solutions (a) across the width at y=0.5H and (b) over the height at x=0.5W, and the NSF, full DSMC and final hybrid temperature solutions (c) across the width at y=0.5H and (d) over the height at x=0.5W.

Compared with the full DSMC solution, the mean error in the NSF velocity solution is 109.23%; Fig. 7.19(a) shows that the hybrid method is able to reduce this to only 0.87% inside 7 iterations. The error in the initial NSF temperature field is 26.73% for this case — this is considerably larger than for the previous non-isothermal test case. Nevertheless, a reduction to only 0.56% is achieved by the hybrid method inside 7 iterations, as indicated in Fig. 7.19(b).

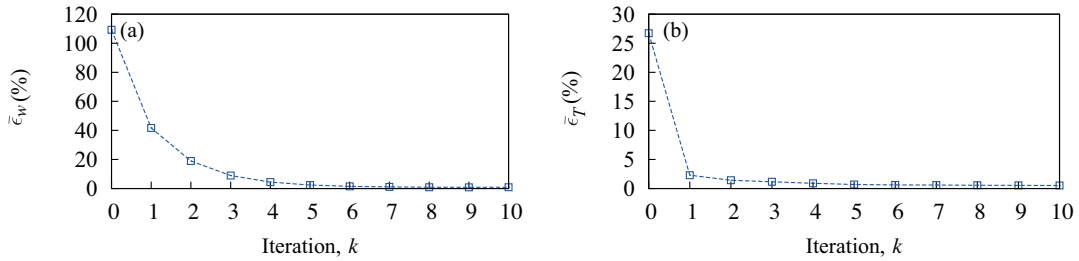

Figure 7.19: Crack flow with a temperature difference over the height: mean error in the hybrid (a) velocity field and (b) temperature field at each iteration; k=0 represents the NSF solution.

The mass flow rate predicted by the full DSMC simulation of this case is $\dot{m}_{\text{Full}}$= 74.41×10^{-11} kg/s. Compared with this value, the error in the NSF prediction (which is again $\dot{m}_{\text{NSF}}$= 5.74×10^{-11} kg/s) is 92.93%. Fortunately, after 7 iterations, the hybrid method estimates $\dot{m}_{\text{Hyb}}$= 74.42×10^{-11} kg/s, equating to an error of only 0.013%.

The high accuracy of the hybrid predictions for this case is a result of the method's ability to approximate the true form of all four of the constitutive correction fields, i.e. those calculated from the full DSMC property fields that are shown in Fig. 7.20. The agreement between the hybrid approximations and these full DSMC corrections is confirmed by the profiles across the width and over the height, shown in Fig. 7.21.

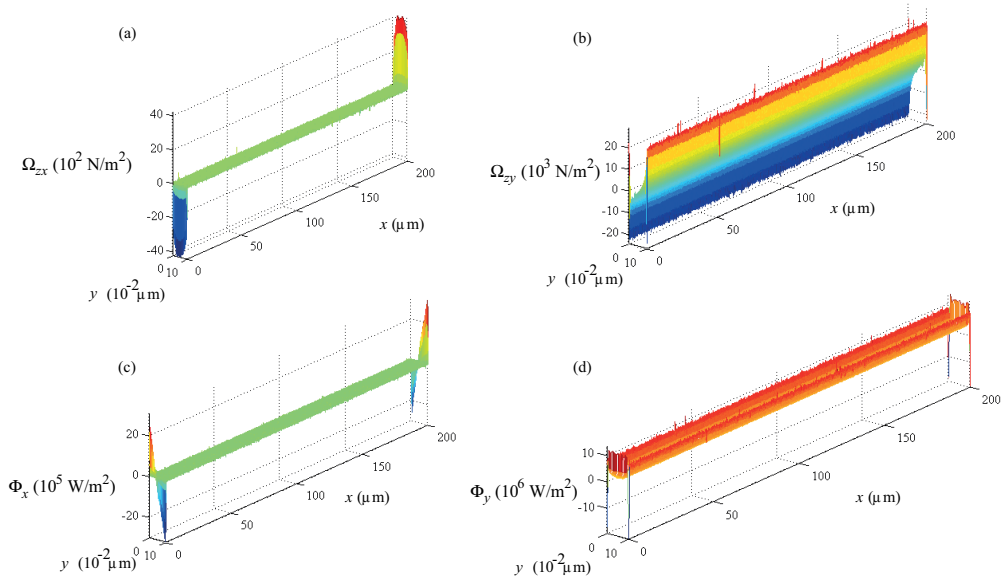

Figure 7.20: Crack flow with a temperature difference over the height: surface plot of the full DSMC (a) Ω_{zx}, (b) Ω_{zy}, (c) Φ_x, and (b) Φ_y correction fields.

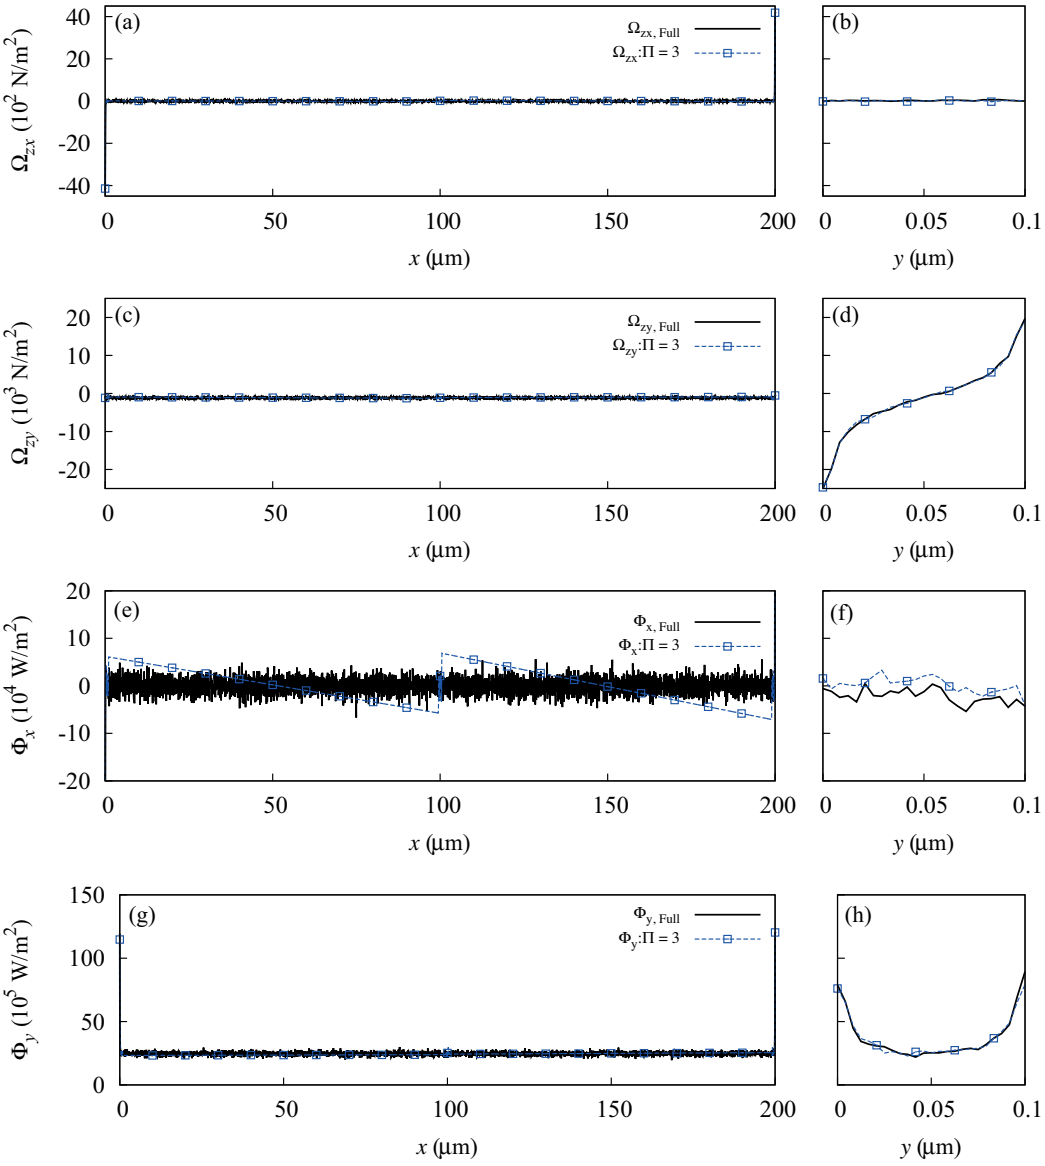

Figure 7.21: Crack flow with a temperature difference over the height: Full DSMC and final shear stress corrections, Ω_{zx} over (a) y=0.5H and (b) x=0.5W, and Ω_{zy} over (c) y=0.5H and (d) x=0.5W. Similarly, the full DSMC and final heat flux corrections, Φ_x over (e) y=0.5H and (f) x=0.5W, and Φ_y over (g) y=0.5H and (h) x=0.5W.

Once again, scatter in the bulk micro element results in a discontinuity in the method's approximation of Φ_x, which can be seen near the centre of the width in Fig. 7.21(c). As for the previous test case, the effects of this discontinuity are not significant.

7.4 Computational savings

Like the test cases of the two previous chapters, each of the cases in this chapter was simulated using the same number of DSMC time steps for the hybrid micro DSMC simulations and the equivalent full DSMC simulation, i.e. $G_{\text{tot,Hyb}} = G_{\text{tot,Full}} = 2\times10^6$. For each case, the computational speed-up S obtained by exploiting spatial scale separation is computed from Eqs. (5.8) and (5.9). With side-wall element extents $W_{\text{left}}=(W_{SR}+W_{RR})_{\text{left}}$ and $W_{\text{right}}=(W_{SR}+W_{RR})_{\text{right}}$ equal to $10\lambda_l$, and bulk element extents $W_{\text{bulk}}=(W_{SR}+2W_{RR})_{\text{bulk}}$ of $15\lambda_l$, Table 7.1 shows that considerable speed-ups are achieved by all of the hybrid simulations performed in this chapter, i.e. $S>1$. Note that, as an additional iteration is required for convergence, the speed-up is slightly less for the non-isothermal case of section 7.3.3.

	Π	W_{left} (μm)	W_{bulk} (μm)	W_{right} (μm)	I	S
isothermal	2	0.98	-	0.98	6	16.84
	3	0.98	1.47	0.98	6	9.95
temperature difference over width	3	1.1	1.45	0.9	6	9.77
temperature difference over height	3	0.98	1.47	0.98	7	8.51

Table 7.1: Computational speed-ups S offered by all hybrid simulations of Chapter 7. The micro element extents W that are stated are those in iteration I and subsequent iterations.

As mentioned previously, the increase in expense that comes from adding a bulk micro element brings no additional accuracy to the isothermal case of section 7.3.1. Although the element configuration used to simulate the non-isothermal cases of sections 7.3.2 and 7.3.3 included a bulk element, two near-wall elements and linear interpolations of the constitutive corrections (and boundary information) may have been adequate for both of these cases: the temperature gradient on the upper and lower walls in section 7.3.2 was linear and so (for the temperature range considered) the transport properties would have varied approximately linearly across the width; in section 7.3.3, the temperature difference over the height produces little variation of the flow properties across the width. Bulk elements could, however, provide essential micro resolution if, for example, there are surface defects, or the cross-section is not perfectly rectangular. A more extreme temperature range may also require bulk micro resolution.

7.5 Summary

The HMM-FWC continuum-DSMC method has been used to simulate gas flowing in a crack-type microchannel that has an aspect-ratio of 2000. Non-isothermal conditions required the method's coupling strategy to deal with the transport of both momentum and heat in 2D. For all of the test cases considered, convergence of the method's algorithm was achieved inside $6-7$ iterations. Like the test cases of the previous two chapters, the results from all cases were compared with equivalent full DSMC solutions. The initial NSF predictions of the velocity fields, the temperature fields, and the mass flow rates were found to have errors in ranges of $98-113\%$, $3-27\%$, and $92-93\%$, respectively. The hybrid method was able to reduce all of these errors to under 1% in all cases and, at the same time, provide substantial speed-ups between $8\times$ and $17\times$ depending (mainly) on the micro element configuration. Importantly, variation in the flow properties across the larger dimension of the cross-section, i.e. the width, was modelled accurately and efficiently — such variation is often neglected in the literature.

The speed-up offered by the hybrid method should increase with increasing cross-sectional aspect-ratio — narrow microscale cracks may have much higher aspect-ratios than that examined in this chapter, for example of $\mathcal{O}(10^5)$ and above. As well as higher aspect-ratios, future work could include the study of this 2D crack problem with higher levels of rarefaction (i.e. larger Kn_{gl}), with non-rectangular cross-sections, or with the presence of surface defects/roughness. More extreme temperatures and larger temperature differences may also require the consideration of momentum transport in the x- and y-directions. All of these scenarios would likely necessitate bulk micro elements, and could perhaps benefit from the use of surface fitting (i.e. 3D curve fitting) techniques to approximate the global constitutive corrections. As mentioned previously, 3D crack-type flows could be studied by implementing the HMM-FWC along the streamwise length of the crack geometry or, if the scale separation in the streamwise direction is high, by combining the HMM-FWC with the IMM. Investigations might then focus on the effects of different streamwise pressure and temperature gradients.

Chapter 8

Conclusions

8.1 Summary

To avoid the computational expense of using full DSMC simulations to model the flow behaviour of non-equilibrium dilute gases, this thesis has focused on the development of a continuum-DSMC hybrid method. The vast majority of existing continuum-DSMC methods are based on a DD framework, and so are effective only for flow problems that can be classed as Type A, i.e. where the micro resolution is needed in localised spatial regions. Heterogeneous frameworks offer an alternative to DD for Type B problems, which require micro resolution everywhere in order to compensate for the failure of the traditional NSF constitutive relations.

In general, continuum-particle hybrid methods achieve computational savings over full particle simulations by exploiting scale separation. However, with the exception of the IMM (which is restricted to flows through very long, thin channels), there has not yet been a heterogeneous continuum-DSMC method that is able to exploit spatial scale separation. A traditional HMM framework has been adopted by a number of continuum-MD hybrid methods for liquid flows; the point-wise coupling approach of these methods is able to exploit highly separated spatial scales. The position and size of the micro elements is, however, restricted by the macro resolution, and this type of framework is not suitable for flows with low or mixed spatial scale separation. Heterogeneous field-wise coupling (HMM-FWC) is able to avoid these restrictions: micro elements of any size can be placed at any location in the flowfield; interpolations across

the distributed micro elements supply the continuum-fluid description with global constitutive and boundary information, meaning that flows with low/mixed spatial scale separation can be simulated. Due to its generality, this HMM-FWC framework provides the basis for the continuum-DSMC method of this thesis.

The original continuum-MD HMM-FWC method was proposed with assumptions of steady, incompressible, and isothermal flow. With the focus remaining on steady flows, the work of this thesis has adapted the HMM-FWC to simulate dilute gas flows: the constraint of the micro elements (i.e. the imposition of the local continuum property fields) has been tailored to a DSMC particle solver. In addition, the coupling strategy has been extended, providing capabilities in compressible and non-isothermal flows. The method has also been made more flexible by enabling the spatial extent of each micro element to adjust dynamically with the local mean free path of the gas.

The versatility of this HMM-FWC method provides the potential to simulate a wide range of flows, including Type A and Type B problems. Its ability to cope with inaccurate boundary and constitutive information has been assessed in this thesis by modelling two 1D validation problems, comparing the accuracy and computational expense with equivalent full DSMC simulations. A micro Fourier flow problem under various temperature and rarefaction conditions was used to test the 1D energy coupling procedure; a high-speed micro Couette flow problem (where viscous heating is present in the gas) with supersonic and hypersonic wall speeds was then able to demonstrate the full coupling algorithm in 1D. Featuring velocity slip/temperature jump and the associated momentum/thermal Knudsen layers, these 1D flow problems could be classed as Type A. However, with the transport property models assumed unknown, constitutive corrections are required everywhere in the flowfield and these problems become Type B.

Convergence of the hybrid algorithm was found to occur very quickly (within 3 or 4 iterations) for most of these 1D validation cases. The accuracy of the converged solutions was found to depend heavily on the arrangement of the micro elements, as this determined how well the method was able to approximate the constitutive correction fields and boundary information. With sufficient micro resolution, the method was able to obtain excellent agreement with the full DSMC solutions. The level of spatial

scale separation was, however, very low in these problems, meaning that the micro elements were required to cover a considerable portion of the flowfield to achieve high accuracy. Unfortunately, this resulted in only moderate computational speed-ups for the Fourier flow test cases, while no speed-ups were obtained for the Couette flow test cases. Nevertheless, the purpose of these 1D cases was simply to validate the method's coupling strategy. Also, for any given flow problem, the micro resolution can be adjusted by the user to achieve the desired balance of accuracy and efficiency.

Following this 1D validation, the method's performance was tested on a larger and more challenging 2D problem. A microchannel with a high-aspect-ratio cross-section was considered as a representative geometry to model the flow of a gas through a microscale crack; leakage from microscale cracks can occur in a variety of engineering systems. The implementation of the method was simplified to 2D by assuming periodicity in the streamwise direction and using an external body forcing to drive the flow. This high-aspect-ratio crack-type geometry can be viewed as a specific class of Type C problem: DD methods are unsuitable as the majority (if not all) of the flow is near-wall; the point-wise coupling approach of the HMM is also unsuitable as the resolution required over the smaller dimension of the cross-section (i.e. the channel height) would result in overlapping of the micro elements. Although the IMM is suitable for Type C flows where high scale separation exists in the streamwise direction, it cannot exploit scale separation over the larger dimension of the channel cross-section (i.e. the width). The HMM-FWC method of this thesis is able to exploit scale separation in the direction of the width by dispersing micro elements across this dimension; at the same time, it is able to deal with the low scale separation over the channel height by setting each element to occupy the entire height.

One isothermal and two non-isothermal test cases were considered based on an aspect ratio of 2000 and a Knudsen number of 1; the results from the hybrid method were again compared with equivalent full DSMC solutions. The isothermal test case was able to demonstrate the momentum coupling procedure in 2D. With temperature differences between the bounding walls of the cross-section, the non-isothermal test cases then demonstrated simultaneous momentum and energy coupling in 2D. For each of the three test cases, the hybrid method's predictions of the property fields and the

mass flow rate were within 1% of the full DSMC solution. Although more iterations (i.e. $6-7$) were required for convergence than were typically required by the 1D test cases of previous chapters, the method was able to achieve substantial computational speed-ups ranging from approximately $8\times$ to $17\times$ depending mainly on the micro element configuration. Even greater speed-ups are expected for higher aspect-ratios.

8.2 Future work

This thesis opens a number of paths for potential future work. For the high-aspect-ratio crack-type flow problem, it would be interesting to consider the following:

- More extreme temperature conditions and/or larger temperature gradients. If the transport properties do not vary linearly over the temperature range of the problem, 'bulk' micro elements would be required. Approximations of the constitutive correction fields might also benefit from the use of surface fitting techniques.

- Geometric variation in the direction of the width. In reality, the geometry of a crack will not be perfectly rectangular, e.g. it could be more open at one end than the other. Surface roughness or even surface defects could also have substantial effects on the flow behaviour. Note that any geometric variation would also likely require bulk micro elements.

- An extension of the method to 3D. The variation of the flow behaviour may not be negligible along the streamwise length of a crack/channel. The HMM-FWC method could be applied along the length to study the effects of streamwise pressure and temperature gradients — this would require a 3D continuum description, constraints of the micro elements in 3D, and 3D interpolations/fitting. Alternatively, if there is high scale separation along the streamwise direction, the IMM could be employed to achieve even greater computational speed-ups (due the streamwise periodicity of the micro elements used in the IMM).

The hybrid method could also be applied to tackle the following:

- Other flow problems with mixed levels of spatial scale separation. This could include the microchannel/reservoirs flow problem that was discussed in Chapter 4.

As mentioned previously, this type of problem could be modelled using the HMM-FWC: small elements distributed along the microchannel length could exploit the higher scale separation in this dimension, while larger micro elements could cope with the lower scale separation in the reservoirs. Again, greater speed-ups might be possible by replacing the HMM-FWC with the IMM along the channel length.

- The flow of polyatomic gases and gas mixtures. Although modelling a monatomic gas was sufficient to validate and test the hybrid method in this thesis, engineering problems involve the flow of gases like air, carbon dioxide, or even hydrogen. The hybrid method has the ability to deal with unknown transport property models in any gas flow. Note that the discussion in Chapter 2 regarding property measurement and control within the DSMC method was based on the assumption of a monatomic gas; while the *dsmcFoamStrath* code is able to deal with the additional complexities in the extraction of the flow properties from polyatomic gas mixtures, the control of particles that possess rotational energy will require further development.

- External high-speed, high-altitude gas flows, for example around hypersonic space access vehicles; physical data is scarce in this area due to the high cost of performing experiments. The hybrid method could be useful in modelling the thin bow shocks that appear at the front of hypersonic vehicles, and the low density wake region behind them. Although the non-equilibrium behaviour of these flows has similarities to that in micro devices, very high temperatures may result in chemical reactions becoming important — the ability to model such reactions would require further investigation and development of the method.

Development of the hybrid method itself could focus on:

- The imposition of a Chapman-Enskog (CE) distribution at the outer boundaries of the micro element relaxation regions. As this distribution incorporates a perturbation from equilibrium, its use may reduce the extents required by the relaxation regions. A CE distribution could be implemented at these boundaries using the technique proposed by Stephani *et al.* [120].

- The selection or definition of a breakdown parameter. Thus far, the initial placement of the micro elements has been based on the author's judgement, and the knowledge that non-equilibrium flow typically occurs close to bounding walls. Existing DD hybrid methods often use a breakdown parameter to identify localised spatial regions that require micro resolution. Although the aim of this hybrid method is to supply a global correction to the continuum description, a breakdown parameter could be useful in determining the placement of the elements; it would be very useful to be able predict where large gradients in the constitutive corrections might occur.

- An extension to unsteady flows. The method could be implemented using the CAI scheme of Lockerby *et al.* [88]. This adaptive time stepping approach is able to deal with mixed/varying degrees of temporal scale separation, and could increase the computational speed-up offered by the method considerably.

- The use of a CFD solver from the OpenFOAM toolbox. Finite difference approximations have been sufficient for the 1D and 2D problems considered in this thesis. Larger and more complex 2D and 3D flow problems may, however, require the use of a CFD solver; *rhoSimpleFoam*, for example, can simulate compressible, steady-state flows.

- The migration of the entire hybrid algorithm to the OpenFOAM toolbox. Currently, the hybrid algorithm is performed by a MATLAB script that creates the micro DSMC simulations, executes both solvers, and performs the information exchange between the solvers. Although a MATLAB script could also be used to couple two separate OpenFOAM solvers (i.e. a CFD solver and *dsmcFoam*), the execution of the entire algorithm could be migrated into OpenFOAM. This package has the benefit of being fully open source, avoiding the need for expensive license fees. Having all of hybrid method's code in the same toolbox and in the same language would also ease future development.

References

[1] R. P. Feynman. There's plenty of room at the bottom. *Talk delivered at the annual meeting of the American Physical Society, Pasadena, CA.*, 1959.

[2] S. U.-S. Choi. Nanofluid technology: current status and future research. *Technical Report. Argonne National Laboratory, Chicago, IL.*, 1999.

[3] H. A. Stone, A. D. Stroock, and A. Ajdari. Engineering flows in small devices: microfluidics toward a lab-on-a-chip. *Annual Review of Fluid Mechanics*, 36:381–411, 2004.

[4] B. Corry. Designing carbon nanotube membranes for efficient water desalination. *Journal of Physical Chemistry B*, 112(5):1427–1434, 2008.

[5] E. B. Arkilic, M. A. Schmidt, and K. S. Breuer. Gaseous slip flow in long microchannels. *Journal of Microelectromechanical Systems*, 6(2):167–178, 1997.

[6] S. Colin. Rarefaction and compressibility effects on steady and transient gas flows in microchannels. *Microfluidics and Nanofluidics*, 1(3):268–279, 2005.

[7] D. A. Lockerby and J. M. Reese. On the modelling of isothermal gas flows at the microscale. *Journal of Fluid Mechanics*, 604:235–261, 2008.

[8] K. C. Lea and S. K. Loyalka. Motion of a sphere in a rarefied gas. *Physics of Fluids*, 25 (9):1550–1557, 1982.

[9] M. D. Allen and O. G. Raabe. Slip correction measurements of spherical solid aerosol particles in an improved millikan apparatus. *Aerosol Science and Technology*, 4(3):269–286, 1985.

[10] B. J. Alder and T. E. Wainwright. Studies in molecular dynamics. I. General method. *Journal of Chemical Physics*, 31(2):459–466, 1959.

[11] G. A. Bird. *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*. Oxford Science Publications, Oxford University Press Inc, New York, 1994.

[12] D. C. Wadsworth and D. A. Erwin. Two-dimensional hybrid continuum/particle approach for rarefied flows. *Proceedings of the 32rd AIAA Plasmadynamics and Lasers Conference, Nashville, TN.*, 1992.

[13] R. Roveda, D. B. Goldstein, and P. L. Varghese. Hybrid Euler/particle approach for continuum/rarefied flows. *Journal of Spacecraft and Rockets*, 35:258–265, 1998.

[14] A. L. Garcia, J. B. Bell, W. Y. Crutchfield, and B. J. Alder. Adaptive mesh and algorithm refinement using Direct Simulation Monte Carlo. *Journal of Computational Physics*, 154 (1):134–155, 1999.

[15] R. Roveda, D. B. Goldstein, and P. L. Varghese. Hybrid Euler/Direct Simulation Monte Carlo calculation of unsteady slit flow. *Journal of Spacecraft and Rockets*, 37:753–760, 2000.

[16] O. Aktas and N. R. Aluru. A combined continuum/DSMC technique for multiscale analysis of microfluidic filters. *Journal of Computational Physics*, 178(2):342–372, 2002.

[17] 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.

[18] Q. Sun, I. D. Boyd, and G. V. Candler. A hybrid continuum/particle approach for modeling subsonic, rarefied gas flows. *Journal of Computational Physics*, 194(1):256–277, 2004.

[19] J.-S. Wu, Y.-Y. Lian, G. Cheng, R. P. Koomullil, and K. C. Tseng. Development and verification of a coupled DSMC–NS scheme using unstructured mesh. *Journal of Computational Physics*, 219(2):579–607, 2006.

[20] T. E. Schwartzentruber and I. D. Boyd. A hybrid particle-continuum method applied to shock waves. *Journal of Computational Physics*, 215(2):402–416, 2006.

[21] G. Abbate, C. R. Kleijn, and B. J. Thijsse. Hybrid continuum/molecular simulations of transient gas flows with rarefaction. *AIAA Journal*, 47(7):1741–1749, 2009.

[22] B. John and M. Damodaran. Computation of head-disk interface gap micro flowfields using DSMC and continuum–atomistic hybrid methods. *International Journal for Numerical Methods in Fluids*, 61(11):1273–1298, 2009.

[23] S.-S. Xu, Z.-N. Wu, Q. Li, and Y.-J. Hong. Hybrid continuum/DSMC computation of rocket mode lightcraft flow in near space with high temperature and rarefaction effect. *Computers and Fluids*, 38(7):1394–1404, 2009.

[24] S. Pantazis and H. Rusche. A hybrid continuum-particle solver for unsteady rarefied gas flows. *Vacuum*, 109:275–283, 2014.

[25] Z. Tang, B. He, and G. Cai. Investigation on a coupled Navier-Stokes–direct simulation Monte Carlo method for the simulation of plume flowfield of a conical nozzle. *International Journal for Numerical Methods in Fluids*, 76(2):95–108, 2014.

[26] W. E and B. Engquist. The heterogeneous multiscale methods. *Communications in Mathematical Sciences*, 1(1):87–132, 2003.

[27] M. K. Borg, D. A. Lockerby, and J. M. Reese. Fluid simulations with atomistic resolution: a hybrid multiscale method with field-wise coupling. *Journal of Computational Physics*, 255:149–165, 2013.

[28] M. K. Borg, D. A. Lockerby, and J. M. Reese. A multiscale method for micro/nano flows of high aspect ratio. *Journal of Computational Physics*, 233:400–413, 2013.

[29] D. A. Kessler, E. S. Oran, and C. R. Kaplan. Towards the development of a multiscale, multiphysics method for the simulation of rarefied gas flows. *Journal of Fluid Mechanics*, 661:262–293, 2010.

[30] A. Patronis, D. A. Lockerby, M. K. Borg, and J. M. Reese. Hybrid continuum–molecular modelling of multiscale internal gas flows. *Journal of Computational Physics*, 255:558–571, 2013.

[31] A. Patronis and D. A. Lockerby. Multiscale simulation of non-isothermal microchannel gas flows. *Journal of Computational Physics*, 270:532–543, 2014.

[32] M. K. Borg, D. A. Lockerby, and J. M. Reese. A hybrid molecular–continuum method for unsteady compressible multiscale flows. *Journal of Fluid Mechanics*, 768:388–414, 2015.

[33] J. C. Maxwell. On stresses in rarefied gases arising from inequalities of temperature. *Philosophical Transactions of the Royal Society Part 1*, 170:231–256, 1879.

[34] M. von Smoluchowski. Üeber wärmeleitung in verdünnten gasen. *Annalen der Physik*, 300(1):101–130, 1898.

[35] N. Dongari. *Micro Gas Flows: Modelling the Dynamics of Knudsen Layers*. PhD thesis, University of Strathclyde, 2012.

[36] C. White. *Benchmarking, Development and Applications of an Open Source DSMC Solver*. PhD thesis, University of Strathclyde, 2013.

[37] M. Gad el Hak. The fluid mechanics of microdevices – The Freeman Scholar Lecture. *Journal of Fluids Engineering*, 121:5–33, 1999.

[38] W. W. Liou and Y. Fang. *Microfluid Mechanics*. McGraw-Hill, Two Penn Plaza, New York, 2006.

[39] J. M. Reese, M. A. Gallis, and D. A. Lockerby. New directions in fluid dynamics: non-equilibrium aerodynamic and microsystem flows. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 361(1813): 2967–2988, 2003.

[40] S. Chapman and T. G. Cowling. *The Mathematical Theory of Non-Uniform Gases*. Cambridge University Press, 1970.

[41] M. S. Shavaliyev. Super-Burnett corrections to the stress tensor and the heat flux in a gas of Maxwellian molecules. *Journal of Applied Mathematics and Mechanics*, 57(3):573–576, 1993.

[42] L. S. García-Colín, R. M. Velasco, and F. J. Uribe. Beyond the Navier-Stokes equations: Burnett hydrodynamics. *Physics Reports*, 465(4):149–189, 2008.

[43] H. Grad. On the kinetic theory of rarefied gases. *Communications on Pure and Applied Mathematics*, 2:331–407, 1949.

[44] H. Struchtrup and M. Torrilhon. Regularization of Grad's 13 moment equations: derivation and linear analysis. *Physics of Fluids*, 15(9):2668–2680, 2003.

[45] X.-J. Gu and D. R. Emerson. A high-order moment approach for capturing non-equilibrium phenomena in the transition regime. *Journal of Fluid Mechanics*, 636:177–216, 2009.

[46] I.-S. Liu and M. A. Rincon. A boundary value problem in extended thermodynamics – one-dimensional steady flows with heat conduction. *Continuum Mechanics and Thermodynamics*, 16(1-2):109–124, 2004.

[47] X.-J. Gu and D. R. Emerson. A computational strategy for the regularized 13 moment equations with enhanced wall-boundary conditions. *Journal of Computational Physics*, 225(1):263–283, 2007.

[48] M. Torrilhon and H. Struchtrup. Boundary conditions for regularized 13-moment-equations for micro-channel-flows. *Journal of Computational Physics*, 227(3):1982–2011, 2008.

[49] P. Taheri, M. Torrilhon, and H. Struchtrup. Couette and Poiseuille microflows: Analytical solutions for regularized 13-moment equations. *Physics of Fluids*, 21(1):017102, 2009.

[50] X.-J. Gu, D. R. Emerson, and G.-H. Tang. Kramers' problem and the Knudsen minimum: a theoretical analysis using a linearized 26-moment approach. *Continuum Mechanics and Thermodynamics*, 21(5):345–360, 2009.

[51] X.-J. Gu and D. R. Emerson. Modeling oscillatory flows in the transition regime using a high-order moment method. *Microfluidics and Nanofluidics*, 10(2):389–401, 2011.

[52] A. Rana, M. Torrilhon, and H. Struchtrup. A robust numerical method for the R13 equations of rarefied gas dynamics: Application to lid driven cavity. *Journal of Computational Physics*, 236:169–186, 2013.

[53] P. L. Bhatnagar, E. P. Gross, and M. Krook. A model for collision processes in gases. i. small amplitude processes in charged and neutral one-component systems. *Physical Review*, 94(3):511–525, 1954.

[54] D. Valougeorgis and S. Naris. Acceleration schemes of the discrete velocity method: Gaseous flows in rectangular microchannels. *Siam Journal on Scientific Computing*, 25 (2):534–552, 2003.

[55] S. Naris and D. Valougeorgis. The driven cavity flow over the whole range of the Knudsen number. *Physics of Fluids*, 17(9):097106, 2005.

[56] S. Pantazis and D. Valougeorgis. Rarefied gas flow through a cylindrical tube due to a small pressure difference. *European Journal of Mechanics-B/Fluids*, 38:114–127, 2013.

[57] L. Mieussens and H. Struchtrup. Numerical comparison of Bhatnagar–Gross–Krook models with proper Prandtl number. *Physics of Fluids*, 16(8):2797–2813, 2004.

[58] A. Frezzotti, G. P. Ghiroldi, and L. Gibelli. Solving the Boltzmann equation on GPUs. *Computer Physics Communications*, 182(12):2445–2453, 2011.

[59] O. Rovenskaya and G. Croce. Numerical investigation of the effect of boundary conditions for a highly rarefied gas flow using the GPU accelerated Boltzmann solver. *Computers and Fluids*, 110:77–87, 2014.

[60] G. A. Bird. Approach to translational equilibrium in a rigid sphere gas. *Physics of Fluids*, 6(10):1518–1519, 1963.

[61] G. A. Bird. The structure of normal shockwaves in a binary gas mixture. *Journal of Fluid Mechanics*, 31(4):657–668, 1968.

[62] G. A. Bird. Breakdown of translational and rotational equilibrium in gaseous expansions. *AIAA Journal*, 8(11):1998–2003, 1970.

[63] C. Borgnakke and P. S. Larsen. Statistical collision model for Monte Carlo simulation of polyatomic gas mixture. *Journal of Computational Physics*, 18(4):405–420, 1975.

[64] G. A. Bird. Perception of numerical methods in rarefied gas dynamics. *Progress in Astronautics and Aeronautics*, 118:211–226, 1989.

[65] G. A. Bird. Direct simulation and the Boltzmann equation. *Physics of Fluids*, 13(11): 2676–2681, 1970.

[66] W. Wagner. A convergence proof for Bird's direct simulation Monte Carlo method for the Boltzmann equation. *Journal of Statistical Physics*, 66(3):1011–1044, 1992.

[67] M. Wang and Z. Li. Simulations for gas flows in microgeometries using the direct simulation Monte Carlo method. *International Journal of Heat and Fluid Flow*, 25(6):975–985, 2004.

[68] R. G. Lord. Application of the Cercignani-Lampis scattering kernel to direct simulation Monte Carlo calculations. *Proceedings of the 17th International Symposium on Rarefied Gas Dynamics*, pages 1427–1433, 1990.

[69] C. C. Su, K. C. Tseng, H. M. Cave, J. S. Wu, Y. Y. Lian, T. C. Kuo, and M. C. Jermy. Implementation of a transient adaptive sub-cell module for the parallel-DSMC code using unstructured grids. *Computers and Fluids*, 39(7):1136–1145, 2010.

[70] A. J. Lofthouse. *Nonequilibrium Hypersonic Aerothermodynamics Using the Direct Simulation Monte Carlo and Navier-Stokes Models.* PhD thesis, University of Michigan, 2008.

[71] OpenFOAM Foundation. http://www.openfoam.org, 2015.

[72] T. J. Scanlon, E. Roohi, C. White, M. Darbandi, and J. M. Reese. An open source, parallel DSMC code for rarefied gas flows in arbitrary geometries. *Computers and Fluids*, 39(10):2078–2089, 2010.

[73] E. Arlemark, G. Markelov, and S. Nedea. Rebuilding of Rothe's nozzle measurements with OpenFOAM software. *Journal of Physics: Conference Series*, 362(1):012040, 2012.

[74] C. White, M. K. Borg, T. J. Scanlon, and J. M. Reese. A DSMC investigation of gas flows in micro-channels with bends. *Computers and Fluids*, 71:261–271, 2013.

[75] M. K. Borg, G. B. Macpherson, and J. M. Reese. Controllers for imposing continuum-to-molecular boundary conditions in arbitrary fluid flow geometries. *Molecular Simulation*, 36(10):745–757, 2010.

[76] M. K. Borg. *Hybrid Molecular-Continuum Modelling of Nano-Scale Flows*. PhD thesis, University of Strathclyde, 2010.

[77] H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. DiNola, and J. R. Haak. Molecular dynamics with coupling to an external bath. *Journal of Chemical Physics*, 81 (8):3684–3690, 1984.

[78] V. I. Kolobov, S. A. Bayyuk, R. R. Arslanbekov, V. V. Aristov, A. A. Frolova, and S. A. Zabelok. Construction of a unified continuum/kinetic solver for aerodynamic problems. *Journal of Spacecraft and Rockets*, 42:598–606, 2005.

[79] V. I. Kolobov, R. R. Arslanbekov, V. V. Aristov, A. A. Frolova, and S. A. Zabelok. Unified solver for rarefied and continuum flows with adaptive mesh and algorithm refinement. *Journal of Computational Physics*, 223(2):589–608, 2007.

[80] K. Xu and J.-C. Huang. A unified gas-kinetic scheme for continuum and rarefied flows. *Journal of Computational Physics*, 229(20):7747–7764, 2010.

[81] S. Tiwari and A. Klar. An adaptive domain decomposition procedure for Boltzmann and Euler equations. *Journal of Computational and Applied Mathematics*, 90(2):223–237, 1998.

[82] J. M. Burt and I. D. Boyd. A low diffusion particle method for simulating compressible inviscid flows. *Journal of Computational Physics*, 227(9):4653–4670, 2008.

[83] J. M. Burt and I. D. Boyd. A hybrid particle approach for continuum and rarefied flow simulation. *Journal of Computational Physics*, 228(2):460–475, 2009.

[84] T. R. Deschenes. *Extension of a Modular Particle-Continuum Method for Nonequilibrium, Hypersonic Flows*. PhD thesis, University of Michigan, 2011.

[85] J. Fan and C. Shen. Statistical simulation of low-speed rarefied gas flows. *Journal of Computational Physics*, 167(2):393–412, 2001.

[86] W.-L. Wang and I. D. Boyd. Hybrid DSMC-CFD simulations of hypersonic flow over sharp and blunted bodies. *Proceedings of the 36th AIAA Thermophysics Conference, Orlando, Fl.*, 2003.

[87] W.-L. Wang and I. D. Boyd. A new energy flux model in the DSMC-IP method for nonequilibrium flows. *Proceedings of the 36th AIAA Thermophysics Conference, Orlando, Fl.*, 2003.

[88] D. A. Lockerby, C. A. Duque-Daza, M. K. Borg, and J. M. Reese. Time-step coupling for hybrid simulations of multiscale flows. *Journal of Computational Physics*, 237:344–365, 2013.

[89] H. S. Wijesinghe, R. D. Hornung, A. L. Garcia, and N. G. Hadjiconstantinou. Three-dimensional hybrid continuum-atomistic simulations for multiscale hydrodynamics. *Journal of Fluids Engineering*, 126(5):768–777, 2004.

[90] N. G. Hadjiconstantinou and A. T. Patera. Heterogeneous atomistic-continuum representations for dense fluid systems. *International Journal of Modern Physics C*, 8(4):967–976, 1997.

[91] Y.-Y. Lian, K.-C. Tseng, Y.-S. Chen, M.-Z. Wu, J.-S. Wu, and G. Cheng. An improved parallelized hybrid DSMC-NS algorithm. *Proceedings of the 26th International Symposium on Rarefied Gas Dynamics, Kyoto, Japan*, 1084(1):341–346, 2008.

[92] T. E. Schwartzentruber. *A Modular Particle-Continuum Numerical Algorithm for Hypersonic Non-equilibrium Flows*. PhD thesis, University of Michigan, 2007.

[93] T. E. Schwartzentruber, L. C. Scalabrin, and I. D. Boyd. A modular particle–continuum numerical method for hypersonic non-equilibrium gas flows. *Journal of Computational Physics*, 225(1):1159–1174, 2007.

[94] T. E. Schwartzentruber, L. C. Scalabrin, and I. D. Boyd. Hybrid particle-continuum simulations of nonequilibrium hypersonic blunt-body flowfields. *Journal of Thermophysics and Heat Transfer*, 22(1):29–37, 2008.

[95] T. E. Schwartzentruber, L. C. Scalabrin, and I. D. Boyd. Multiscale particle-continuum simulations of hypersonic flow over a planetary probe. *Journal of Spacecraft and Rockets*, 45(6):1196–1206, 2008.

[96] T. E. Schwartzentruber, L. C. Scalabrin, and I. D. Boyd. Hybrid particle-continuum simulations of hypersonic flow over a hollow-cylinder-flare geometry. *AIAA Journal*, 46 (8):2086–2095, 2008.

[97] T. R. Deschenes, T. D. Holman, and I. D. Boyd. Parallelization of modular particle-continuum method for hypersonic, near equilibrium flows. *Proceedings of the 48th AIAA Aerospace Sciences Meeting, Orlando, Fl*, 2010.

[98] T. R. Deschenes, T. D. Holman, and I. D. Boyd. Analysis of internal energy transfer within a modular particle-continuum method. *Proceedings of the 47th AIAA Aerospace Sciences Meeting, Orlando, Fl*, 2009.

[99] T. R. Deschenes, T. D. Holman, and I. D. Boyd. Effects of rotational energy relaxation in a modular particle-continuum method. *Journal of Thermophysics and Heat Transfer*, 25(2):218–227, 2011.

[100] Sean T. O'Connell and Peter A. Thompson. Molecular dynamics - continuum hybrid computations: A tool for studying complex fluid flows. *Physical Review E*, 52:R5792–R5795, 1995.

[101] R. Delgado-Buscalioni and P. V. Coveney. Continuum-particle hybrid coupling for mass, momentum, and energy transfers in unsteady fluid flow. *Physical Review E*, 67:046704, 2003.

[102] E. G. Flekkøy, G. Wagner, and J. Feder. Hybrid model for combined particle and continuum dynamics. *Europhysics Letters*, 52(3):271, 2000.

[103] A. L. Garcia and B. J. Alder. Generation of the Chapman-Enskog distribution. *Journal of Computational Physics*, 140(1):66–70, 1998.

[104] D. B. Hash and H. A. Hassan. A decoupled DSMC/Navier-Stokes analysis of a transitional flow experiment. *Proceedings of the 34th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV*, 353, 1996.

[105] W.-L. Wang and I. D. Boyd. Predicting continuum breakdown in hypersonic viscous flows. *Physics of Fluids*, 15:91–100, 2003.

[106] H. S. Tsien. Superaerodynamics, mechanics of rarefied gases. *Journal of the Aeronautical Sciences*, 13:653–654, 1946.

[107] S. Tiwari. Coupling of the Boltzmann and Euler equations with automatic domain decomposition. *Journal of Computational Physics*, 144(2):710–726, 1998.

[108] D. A. Lockerby, J. M. Reese, and H. Struchtrup. Switching criteria for hybrid rarefied gas flow solvers. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 465(2105):1581–1598, 2009.

[109] I. D. Boyd. Predicting breakdown of the continuum equations under rarefied flow conditions. *Proceedings of the 23rd International Symposium on Rarefied Gas Dynamics, Wistler, Canada*, 663:899–906, 2003.

[110] J. Meng, N. Dongari, J. M. Reese, and Y. Zhang. Breakdown parameter for kinetic modeling of multiscale gas flows. *Physical Review E*, 89(6):063305, 2014.

[111] W. E, W. Ren, and E. Vanden-Eijnden. A general strategy for designing seamless multiscale methods. *Journal of Computational Physics*, 228(15):5437–5453, 2009.

[112] W. Ren and W. E. Heterogeneous multiscale method for the modeling of complex fluids and micro-fluidics. *Journal of Computational Physics*, 204(1):1–26, 2005.

[113] S. Yasuda and R. Yamamoto. A model for hybrid simulations of molecular dynamics and computational fluid dynamics. *Physics of Fluids*, 20(11):113101, 2008.

[114] N. Asproulis, M. Kalweit, and D. Drikakis. A hybrid molecular continuum method using point wise coupling. *Advances in Engineering Software*, 46(1):85–92, 2012.

[115] S. Yasuda and R. Yamamoto. Synchronized Molecular-Dynamics simulation via macroscopic heat and momentum transfer: An application to polymer lubrication. *Physical Review X*, 4(4):041011, 2014.

[116] K. Ritos. *Water Flow at the Nanoscale: a Computational Molecular and Fluid Dynamics Investigation*. PhD thesis, University of Strathclyde, 2014.

[117] M. K. Borg, D. A. Lockerby, and J. M. Reese. A hybrid molecular-continuum simulation method for incompressible flows in micro/nanofluidic networks. *Microfluidics and Nanofluidics*, 15(4):541–557, 2013.

[118] D. Stephenson, D. A. Lockerby, M. K. Borg, and J. M. Reese. Multiscale simulation of nanofluidic networks of arbitrary complexity. *Microfluidics and Nanofluidics*, 18:841–858, 2015.

[119] N. G. Hadjiconstantinou Massachusetts Institute of Technology. http://web.mit.edu/ngh-group/research3-variance.html, 2014.

[120] K. A. Stephani, D. B. Goldstein, and P. L. Varghese. A non-equilibrium surface reservoir approach for hybrid DSMC/Navier–Stokes particle generation. *Journal of Computational Physics*, 232(1):468–481, 2013.

[121] B. A. Younglove and H. J. M. Hanley. The viscosity and thermal conductivity coefficients of gaseous and liquid argon. *Journal of Physical and Chemical Reference Data*, 15:1323–1337, 1986.

[122] Mathworks. http://uk.mathworks.com/help/matlab/ref/polyfit.html, 2014.

[123] Mathworks. http://uk.mathworks.com/matlabcentral/fileexchange/13812-splinefit, 2014.

[124] C. Hong, Y. Asako, and J.-H. Lee. Estimation of leak flow rates through narrow cracks. *Journal of Pressure Vessel Technology*, 131(5):051405, 2009.

[125] X. Chong. Subsonic choked flow in the microchannel. *Physics of Fluids*, 18:127104, 2006.

[126] E. Roohi, M. Darbandi, and V. Mirjalili. Direct simulation Monte Carlo solution of subsonic flow through micro/nanoscale channels. *Journal of Heat Transfer*, 131(9):092402, 2009.