

POLITECNICO DI TORINO


Master's Degree in Aerospace Engineering

Technische Universität Dresden

Space Transportation Systems Research Group



**Politecnico
di Torino**



**Technische
Universität
Dresden**

Master Degree Thesis

Breakdown Criteria for DSMC-CFD Coupling in Rarefied Plume-Surface Interactions

Supervisors:

Prof. Andrea FERRERO

Dr.-Ing. Christian BACH

Dipl.-Ing Jannis PETERSEN

Candidate:

Natia SAKANDELIDZE

March/April 2026

Acknowledgements

დიდი მადლობა ჩემს ბებიას და ბაბუს. რომლებმაც ჩემი ცხოვრების საფუძველი ჩაყარეს. თუ დღეს ის ადამიანი ვარ, ვინც ვარ, ეს მათი დამსახურებაა.

Vorrei ringraziare mia mamma per avermi sempre sostenuto nelle mie decisioni e ambizioni, e per tutto quello che ha fatto per portarmi dove sono oggi.

Grazie a Gino che ha sempre fatto ciò che era meglio per me.

Noemi e Martina, sono davvero felice di essere uscita quel giorno per incontrare i nuovi compagni di corso, perché ho conosciuto una persona meravigliosa. Gli anni passati a vivere con voi sono tra i ricordi più belli che ho, e li porterò sempre con me. Vi voglio tanto bene e, qualsiasi cosa succeda in futuro, una parte di me resterà sempre con voi.

Grazie a Davide, da quando ti ho conosciuto sono cambiata molto. Siamo cresciuti insieme e mi hai sempre incoraggiata, anche nei momenti in cui facevo fatica a credere in me stessa.

Grazie al mio gruppo di amici a Torino, sia per il divertimento e i ricordi bellissimi, sia per tutte le esperienze condivise. È vero che negli ultimi tempi di questo percorso ho avuto meno occasioni per stare con voi, ma mi renderebbe davvero felice potervi chiamare amici anche tra tanti anni.

Vorrei ringraziare il professor Ferrero per avermi seguito durante questa tesi, essendo per me un punto di riferimento importante e dandomi fiducia in quello che facevo.

Thanks to Jannis and the entire research department in Dresden for hosting me and supporting me during this time, and to Daniele and Michele for the coffee breaks, the laughs, and for making everything feel lighter.

Thanks to my flatmates for the amazingly weird conversations, for the shared dinners, and for making me feel at home even in a different country. I understood Germany a little bit better thanks to you guys.

Abstract

The establishment of a sustainable human presence on the Moon requires reliable and efficient surface operations. A critical challenge for lunar landers is the interaction between rocket exhaust plumes and the lunar regolith, which can mobilize high-velocity dust particles and pose risks to equipment and infrastructure.

This thesis investigates the modelling of plume-surface interaction (PSI) under near-vacuum conditions, where the expanding exhaust transitions from dense, continuum flow near the nozzle to rarefied, non-equilibrium conditions farther away. Classical Navier-Stokes-based Computational Fluid Dynamics (CFD) fails in the rarefied region, while particle-based Direct Simulation Monte Carlo (DSMC) methods are computationally expensive in dense regions.

The focus of this work is the evaluation of continuum breakdown parameters for identifying the onset of rarefied effects. CFD simulations are performed for multiple plume stand-off heights, and the results are compared with reference DSMC simulations to evaluate the accuracy of different breakdown parameters. By analysing key flow variables, the study assesses where CFD and DSMC predictions diverge and how well the selected breakdown parameters predict the continuum-to-rarefied transition.

The insights gained contribute to the design and assessment of in-situ manufactured lunar landing platforms and inform future hybrid simulation studies for rarefied lunar and planetary environments.

Contents

1	Introduction	5
2	Theoretical background	9
2.1	The OpenFOAM Framework	9
2.2	Computational Fluid Dynamics	10
2.2.1	Pressure-based and density-based solvers	10
2.2.2	Finite Volume Method	11
2.2.3	Spatial and Temporal Discretisation	12
2.2.4	Courant-Friedrichs-Lewy Condition	12
2.2.5	Turbulence Modelling	13
2.2.6	rhoCentralFoam	14
2.3	DSMC: Direct simulation Monte Carlo	15
2.3.1	Algorithmic Procedure	18
2.3.2	Collision models	20
2.4	Hybrid CFD-DSMC	22
2.4.1	Coupling Strategies and Information Exchange in DDM	24
2.4.2	Breakdown parameters	27
3	System model	35
3.1	Numerical Setup	35
3.2	Boundary Conditions	37
3.3	Turbulence Modelling	37
3.3.1	Inlet Turbulence Conditions	37
3.3.2	Wall Treatment	38
3.3.3	Outlet Treatment	38
3.4	Mesh Design	38
3.4.1	Domain Geometry	39
3.4.2	Mesh generation with ICEM CFD	39
3.4.3	Estimation of the First Boundary-Layer Height	41
3.5	Thermophysical Properties of Nitrogen Gas	44
3.5.1	Equation of State and Thermodynamic Model	44
3.5.2	Transport Properties: Sutherland Law Coefficients for Nitrogen	45
3.5.3	Summary of Properties	46
3.5.4	Model Simplification Reasons	46
3.6	DSMC Simulation Setup and Coupling with CFD Nozzle Flow	46
3.6.1	Variation of DSMC Parameters with Stand-Off Height	47
3.6.2	Gas Model and Molecular Properties	48
3.6.3	Collision Model	48
3.6.4	Boundary Conditions	49

3.6.5	Particle Representation and Sampling	49
3.6.6	Coupling Strategy and Limitations	49
4	Numerical evaluation	51
4.1	Grid study	51
4.2	CFD results	54
4.2.1	Observation on Total Temperature Behaviour in CFD Simulations	58
4.2.2	Effect of Boundary Conditions on Pressure and Density Fields	60
4.3	DSMC results	60
4.4	Centreline Comparison Between CFD and DSMC	63
4.5	Ground Pressure Comparison with Experimental Data	66
4.6	Mean Free Path Estimation	68
4.7	Knudsen Number	71
4.8	Bird's Parameter	73
4.9	Gradient Local Length Knudsen number	75
4.10	Deborah number	77
5	Conclusions	79
	Bibliography	81
A	Interpolating fields using MapFields	83
B	Turbulence Boundary Conditions	87
B.1	Full 0/k files	87
B.2	Full 0/omega file	88
C	CFD and DSMC results and Breakdown contours	91

Chapter 1

Introduction

The establishment of a sustainable human presence on the Moon requires the development of reliable, efficient, and robust surface operations. One of the critical challenges associated with lunar missions is the interaction between rocket exhaust plumes and the lunar regolith during landing and take-off manoeuvres. In the absence of an atmosphere, plume-surface interaction (PSI) generates high-velocity jets that directly impinge on the loose, fine-grained lunar soil. This process can mobilize and accelerate regolith particles to considerable velocities, creating risks for landers, surrounding infrastructure, optical equipment, and scientific instruments. Historical lunar missions, from the Apollo program to recent robotic landers, have highlighted the severity of surface erosion and dust transport, demonstrating the need for dedicated mitigation strategies.

In this context, the *LUNAR ISLANDS* project investigates novel concepts for in-situ manufacturing of lunar landing platforms. Within this activity an international consortium led by TU Dresden is assessing the effectiveness of in-situ manufactured structures on the containment of high-velocity particles that are generated by the plume-surface interaction of landers. This involves the manufacturing and characterisation of samples from lunar regolith simulants, experimental investigations in the form of cold gas and hot fire tests as well as numerical simulations of the interaction of exhaust plumes and lunar regolith. During this project, TU Dresden as the prime is executing a central part of the numerical investigations of the continuum flow and the peripheral region of the exhaust jet.

The numerical simulations performed at TU Dresden provided the foundation for this thesis and motivated the investigation of several questions that arose during the computational work. In particular, modelling PSI in lunar environments introduces unique difficulties. The expanding rocket plume transitions rapidly from high-pressure, high-density conditions near the nozzle to extremely low-density, near-vacuum conditions farther away. This results in flow regimes spanning continuum to rarefied gas dynamics within a single domain. Classical computational fluid dynamics (CFD), based on the Navier-Stokes equations, becomes invalid in regions where the continuum assumption breaks down. Conversely, particle-based methods such as the Direct Simulation Monte Carlo (DSMC) technique accurately capture rarefied effects but are computationally demanding when applied to dense flow regions, computational demands of DSMC scale with Kn^{-4} [13]. A coupled CFD-DSMC approach therefore offers an efficient compromise: CFD resolves the continuum region near the nozzle and along the jet core, while DSMC is used in the rarefied regions where molecular non-equilibrium effects dominate. This coupling of solvers enables the treatment of a wide range of flow problems, including, but not limited to, unconventional gas reservoirs, atmospheric re-entry flows, micro- and nano-scale devices,

stratospheric aerosols, high-pressure shock tubes, sonoluminescence, and liquid-vapour interface phenomena. Accordingly, the numerical simulations presented in this work adopt such a hybrid approach.

A central challenge in constructing such hybrid models is determining where the transition from continuum to rarefied flow occurs. In most practical flows, the transition from continuum to rarefied conditions occurs gradually rather than abruptly, and no single physical quantity provides an exact demarcation. Various continuum breakdown parameters and threshold values have been proposed in the literature, including Knudsen number-based metrics, local gradients, mean free path considerations, or non-equilibrium indicators. Yet, their performance and suitability can depend strongly on the specific flow field, geometry, and physical conditions. The accuracy of the breakdown interface directly affects model fidelity and computational cost. In addition, steep density and velocity gradients typical of supersonic jet expansions and other multi-regime flows can induce non-equilibrium effects even in regions where the local Knudsen number is moderate, further complicating the determination of the interface. A systematic comparison of different breakdown criteria is therefore necessary to support reliable hybrid simulations within such studies.

This thesis focuses on identifying the continuum-rarefied transition in plume-surface flows and on establishing a comprehensive understanding of the parameters involved. Specifically, how different continuum breakdown parameters and their associated threshold values influence the placement of the CFD-DSMC interface. To investigate this, several breakdown criteria are evaluated.

The work is structured around the following contributions:

- Development of a CFD-based PSI simulation framework in OpenFOAM, incorporating geometry, flow configuration, and relevant physics for supersonic plume expansion in near-vacuum conditions.
- Selection and justification of appropriate boundary conditions, guided by experimental configurations within the LUNAR ISLANDS program and supporting theoretical considerations.
- In-depth study of continuum breakdown parameters, including the underlying physical assumptions and recommended threshold values.
- Execution of CFD simulations of the flow for multiple plume stand-off heights to capture different rarefaction levels within the domain.
- Identification of the spatial location of the breakdown interface for at least three different breakdown criteria.
- Quantitative comparison of CFD predictions with data from a mostly DSMC (interface placed at nozzle throat) simulation provided by the project supervisor.
- Comprehensive documentation, analysis, and discussion of the numerical setup, modelling decisions, and resulting findings, as presented within this thesis.

The focus is on key flow variables such as density, pressure, temperature and velocity, so as to analyse discrepancies in and out of the continuum region. Through this analysis, the thesis contributes to the broader effort of advancing accurate and efficient numerical tools for plume-surface interaction studies in rarefied environments. The insights gained support the design and assessment of in-situ manufactured lunar landing platforms and

provide guidance for future hybrid simulation frameworks relevant to lunar exploration and beyond.

The main contributions of this thesis are:

- Development of a numerical setup in OpenFOAM capable of simulating plume expansion into vacuum and generating CFD data for hybrid coupling.
- Implementation of a hybrid CFD–DSMC workflow that enables direct comparison between continuum-based and particle-based methods for multi-regime flows.
- Systematic evaluation of at least three established continuum breakdown parameters, including an assessment of their recommended threshold values under identical flow conditions.
- Quantitative comparison of CFD, hybrid CFD-DSMC, and full DSMC solutions at multiple standoff heights, focusing on density, pressure, velocity, and temperature distributions.
- Identification of the spatial regions where each breakdown parameter predicts loss of continuum validity and analysis of how these predictions correlate with deviations between CFD and DSMC results.
- Documentation of practical considerations for solver setup, mesh design, boundary conditions, and interface placement that influence hybrid simulation accuracy.
- Recommendations for selecting continuum breakdown criteria in future plume-surface interaction studies or any application involving transitions between continuum and rarefied regimes.

The structure of this thesis is as follows. Chapter 2 provides the theoretical background required for the subsequent analysis, including an overview of continuum and rarefied gas dynamics, the governing equations of CFD and DSMC methods, and the established continuum breakdown parameters considered in this work.

Chapter 3 describes the numerical methodology, detailing the CFD setup and DSMC specifics in OpenFOAM.

Chapter 4 presents the simulation results for the investigated stand-off heights, including the CFD and the DSMC fields, interface locations predicted by the various breakdown criteria. A comparative analysis of key flow quantities is carried out to assess the performance of each breakdown parameter.

Chapter 5 discusses the implications of the findings, evaluates the robustness of the criteria and highlights practical considerations for hybrid simulations of multi-regime flows and provides suggestions for future work. Computational resources provided by HPC@PoliTO (www.hpc.polito.it)

Chapter 2

Theoretical background

2.1 The OpenFOAM Framework

OpenFOAM (Open Field Operation and Manipulation) is an open-source CFD software in which the codes are written in C++. It provides a flexible environment in which users can access, modify, and extend numerical solvers and physical models. OpenFOAM implements the Finite Volume Method on arbitrary polyhedral meshes and supports both structured and unstructured geometries, making it suitable for complex engineering configurations.

One of the distinguishing features of OpenFOAM is its modular design. Each simulation is structured within a case directory, typically composed of three main folders:

- `0/` for initial and boundary conditions,
- `constant/` for physical and geometric definitions,
- `system/` for solver and numerical control parameters.

This hierarchy contributes to both transparency and reproducibility of simulations. The framework also employs a runtime selection mechanism, enabling users to switch between different discretisation schemes, turbulence models, and numerical solvers without recompilation.

OpenFOAM offers a wide range of solvers covering incompressible and compressible flows, multiphase systems, combustion, and heat transfer.

For compressible flows, algorithms such as `rhoPisoFoam` and `rhoCentralFoam` are frequently used. These solvers combine pressure-velocity coupling algorithms (PISO, SIMPLE, or hybrid PIMPLE) with flux-limiting schemes to handle strong shocks and discontinuities, which are particularly relevant for supersonic and underexpanded jets.

The selection of OpenFOAM for this work is motivated by several considerations. First, its open-source nature allows for complete transparency of the numerical implementation and direct access to solver routines, which is often restricted in commercial packages. Second, OpenFOAM's native support for compressible, transient, and turbulent flows makes it well-suited for the analysis of jet expansion and shock formation phenomena. Third, its scalability and parallelisation capabilities enable simulations of large computational domains with reasonable efficiency on multi-core systems. Finally, the extensive academic and industrial community around OpenFOAM provides validated cases and best-practice guidelines that support methodological robustness and reproducibility.

2.2 Computational Fluid Dynamics

Computational Fluid Dynamics (CFD) is a branch of fluid mechanics that employs numerical analysis and algorithms to solve and predict problems involving fluid flow, heat transfer, and related physical phenomena. Instead of relying solely on experimental or analytical approaches, CFD allows for the approximation of the governing conservation equations of mass, momentum and energy on a discrete computational grid [25].

The fundamental equations solved in CFD are the Navier-Stokes equations, which describe the motion of a viscous, Newtonian fluid. The equations to be solved are expressed by the following system of differential equations, which express the physical laws of conservation of mass (continuity), conservation of momentum and conservation of energy:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \bar{u}) = 0, \quad (2.1)$$

$$\frac{\partial(\rho \bar{u})}{\partial t} + \nabla \cdot (\rho \bar{u} \bar{u}) = \nabla \cdot \bar{\bar{\sigma}} + \rho \bar{f}_e, \quad (2.2)$$

$$\frac{\partial \rho E}{\partial t} + \nabla \cdot (\rho \bar{u} E) = \nabla \cdot (\bar{\sigma} \bar{u}) + \rho \bar{f}_e \cdot \bar{u} - \Delta \cdot \bar{q}_T \quad (2.3)$$

where ρ is the fluid density, $\bar{u}$ the velocity vector, p the pressure, $\bar{f}_e$ volume forces per unit of mass, $\bar{q}_T$ heat flow and $\bar{\sigma} = -p\bar{n} + \bar{\tau}$ stress tensor with $\bar{\tau}$ viscosity term. These equations are, in general, non-linear and coupled, which makes analytical solutions possible only for highly idealised flows.

2.2.1 Pressure-based and density-based solvers

In CFD, the numerical solution of the Navier-Stokes equations can be approached using either density-based or pressure-based solvers. Density-based solvers, also known as coupled solvers, compute the density, momentum and energy equations simultaneously and are particularly suitable for compressible, high-speed flows where strong density variations, shock waves or expansion fans are present. Pressure-based solvers, also referred to as the segregated solvers, on the other hand, typically solve a pressure-velocity coupling and are more commonly used for incompressible or low-Mach-number flows.

Pressure-based solvers

In pressure-based solvers, the mass and momentum equations are solved in a segregated, uncoupled manner. In other words, the governing equations are treated sequentially to obtain the flow variables. Since these solvers are typically applied to incompressible flows, where the density ρ is constant, the primary variables to be computed are the three velocity components u_i and the pressure p .

For simplicity one dimensional incompressible flow is considered below:

$$\frac{\partial u}{\partial x} = 0 \quad (2.4)$$

$$\frac{\partial(\rho u)}{\partial t} + \frac{\partial}{\partial x}(\rho u u) = -\frac{\partial p}{\partial x} + \frac{\partial}{\partial x}\left(\mu \frac{\partial u}{\partial x}\right) \quad (2.5)$$

Since there is no explicit governing equation for the pressure field in incompressible flows, pressure-based methods were developed to enforce mass conservation by coupling the pressure and velocity fields. Some of the most popular pressure-based algorithms include

SIMPLE, SIMPLEC and PISO. The SIMPLE (Semi-Implicit Method for Pressure-Linked Equations) algorithm has seen many variants over the years, making it one of the most commonly studied algorithms in CFD. In it, an initial estimate of the velocity and pressure fields is specified, and boundary conditions are applied. The discretized momentum equations are then solved to obtain an intermediate (or predicted) velocity field. Using these predicted velocities, the mass fluxes through the cell faces are computed by interpolation. A pressure-correction equation is subsequently derived and solved to obtain the pressure correction values within each cell. These corrections are used to update the pressure, velocity, and mass flux fields to enforce continuity. Additional transport equations, such as those for turbulence quantities or energy, are then solved as required. The entire cycle is repeated until convergence is achieved. The main weakness of this algorithm is the slow convergence, due to the estimated nature of the pressure correction equation.

Density-based solvers

In density-based algorithms, which typically employ Riemann solvers, the governing equations of continuity, momentum, and energy are advanced in time using either explicit or implicit time-marching schemes. These equations are solved simultaneously in a fully coupled manner, while the transport equations for turbulence quantities are generally solved in a segregated fashion outside the coupled solver. For simplicity, the one-dimensional Euler equations for compressible flow are shown below:

$$\frac{\partial \rho}{\partial t} + \frac{\partial(\rho u)}{\partial x} = 0, \quad (2.6)$$

$$\frac{\partial(\rho u)}{\partial t} + \frac{\partial}{\partial x}(\rho u^2 + p) = 0, \quad (2.7)$$

$$\frac{\partial(\rho E)}{\partial t} + \frac{\partial}{\partial x}[u(p + \rho E)] = 0. \quad (2.8)$$

Both pressure-based and density-based solvers obtain the velocity field from the momentum equation. The key difference lies in how the pressure field is determined. In a pressure-based method, a pressure (or pressure-correction) equation is constructed by manipulating the continuity and momentum equations to enforce mass conservation. In contrast, a density-based solver first computes the density field directly from the continuity equation and then evaluates the pressure using an appropriate equation of state.

2.2.2 Finite Volume Method

The governing equations introduced 2.1, 2.2 and 2.3 are solved numerically using the finite volume method (FVM). In this approach, the computational domain is discretised into a finite number of non-overlapping control volumes (cells) and the integral form of the conservation equations is applied to each of them.

For a generic conserved quantity $\mathbf{U}$, the integral form of the conservation law over a control volume V with boundary surface S can be written as

$$\frac{\partial}{\partial t} \int_V \mathbf{U} dV + \oint_S \mathbf{F}(\mathbf{U}) \cdot \mathbf{n} dS = \int_V \mathbf{S} dV, \quad (2.9)$$

where $\mathbf{F}(\mathbf{U})$ represents the flux tensor, $\mathbf{n}$ is the outward-pointing unit normal vector to the surface and $\mathbf{S}$ denotes source terms.

The key idea of the finite volume method is that conservation is enforced locally over each control volume. The surface integral term is approximated as the sum of fluxes across

the faces of each cell, which ensures strict conservation of mass, momentum and energy across the computational domain.

In practice, the variables are stored at the cell centres, while fluxes are evaluated at the cell faces. This requires interpolation or reconstruction procedures to obtain face values from neighbouring cells. The accuracy and stability of the numerical solution therefore depend strongly on the choice of discretisation schemes used for these fluxes.

The finite volume formulation is particularly well suited for compressible flows and problems involving discontinuities, such as shock waves, since it naturally preserves conservation properties even in the presence of strong gradients.

2.2.3 Spatial and Temporal Discretisation

Following the finite volume formulation, the governing equations must be discretised both in space and time in order to obtain a system of algebraic equations that can be solved numerically.

Spatial Discretisation

In the finite volume framework, spatial discretisation is achieved by approximating the fluxes across the faces of each control volume. Since the primary variables are stored at cell centres, suitable interpolation or reconstruction schemes are required to evaluate their values at the cell faces.

Different spatial discretisation schemes can be employed, each characterised by a different balance between accuracy and numerical stability. Central differencing schemes provide second-order accuracy but may lead to numerical oscillations in the presence of strong gradients or discontinuities. Upwind schemes, on the other hand, introduce numerical dissipation by biasing the interpolation in the direction of the flow, thereby improving stability at the cost of reduced accuracy.

In compressible flow simulations, particularly those involving shock waves, it is essential to adopt schemes that ensure both stability and sharp resolution of discontinuities. For this reason, hybrid approaches are often used, combining central and upwind features.

Temporal Discretisation

Temporal discretisation consists in approximating the time derivative of the conserved variables. This can be achieved using either explicit or implicit time integration schemes.

In explicit schemes, the solution at the new time level is computed directly from known quantities at the current time level. These methods are straightforward to implement and computationally efficient per time step, but they are conditionally stable and require the time step to satisfy a stability constraint. In contrast, implicit schemes involve the solution of a system of equations at each time step, which increases computational cost but allows for larger time steps and improved stability.

The choice of spatial and temporal discretisation schemes plays a crucial role in determining the accuracy, stability, and computational cost of the simulation, and must therefore be selected consistently with the physical characteristics of the problem under consideration.

2.2.4 Courant-Friedrichs-Lewy Condition

The stability of explicit time integration schemes is governed by the Courant-Friedrichs-Lewy (CFL) condition, which imposes a restriction on the maximum allowable time step.

This condition ensures that the numerical domain of dependence encompasses the physical domain of dependence of the governing equations.

The CFL number is defined as

$$\text{CFL} = \frac{u \Delta t}{\Delta x}, \quad (2.10)$$

where u is a characteristic flow velocity, Δt is the time step and Δx is a representative cell size. For multidimensional and compressible flows, the definition is extended to account for all relevant wave speeds, including the speed of sound, leading to a more general form

$$\text{CFL} = \frac{(|\mathbf{u}| + c) \Delta t}{\Delta x}, \quad (2.11)$$

where c is the local speed of sound.

Physically, the CFL condition requires that information carried by the flow does not travel more than one computational cell during a single time step. If this condition is violated, the numerical scheme may become unstable, leading to non-physical oscillations or divergence of the solution.

For explicit solvers such as `rhoCentralFoam`, the time step must be chosen such that the CFL number remains below a critical value, typically of order unity, throughout the computational domain. In practice, a more restrictive value is often adopted to ensure stability in regions with strong gradients, such as shock waves.

As a consequence, the time step is dynamically adjusted during the simulation based on the local flow conditions and mesh resolution. This requirement can significantly increase the computational cost, particularly for high-speed flows or finely resolved meshes, but it is essential for maintaining the stability and accuracy of the solution.

2.2.5 Turbulence Modelling

In many practical engineering applications, fluid flows are turbulent, characterised by chaotic and multi-scale fluctuations in velocity and pressure. Direct numerical simulation (DNS) of such flows, in which all turbulent scales are fully resolved, is computationally prohibitive for most realistic problems. As a result, modelling approaches are required.

Reynolds-Averaged Navier-Stokes (RANS) Equations

A widely used approach is based on the Reynolds-Averaged Navier-Stokes (RANS) equations, in which the instantaneous flow variables are decomposed into mean and fluctuating components. For a generic variable ϕ , this decomposition is written as

$$\phi = \bar{\phi} + \phi', \quad (2.12)$$

where $\bar{\phi}$ denotes the time-averaged component and ϕ' represents the fluctuating part.

For compressible turbulent flows, the averaging procedure is typically performed using Favre (density-weighted) averaging. In this approach, a generic flow variable ϕ is decomposed as

$$\phi = \tilde{\phi} + \phi'', \quad (2.13)$$

where $\tilde{\phi}$ is the Favre-averaged quantity, defined as

$$\tilde{\phi} = \frac{\overline{\rho\phi}}{\bar{\rho}}, \quad (2.14)$$

and ϕ'' represents the corresponding fluctuation.

Density and pressure are decomposed using Reynolds, while other variables are done with Favre. Applying this decomposition to the Navier-Stokes equations and performing the averaging procedure leads to the compressible RANS equations. The momentum equation can be written as

$$\frac{\partial(\bar{\rho}\tilde{u}_i)}{\partial t} + \frac{\partial(\bar{\rho}\tilde{u}_i\tilde{u}_j)}{\partial x_j} = -\frac{\partial\bar{p}}{\partial x_i} + \frac{\partial}{\partial x_j} \left(\bar{\tau}_{ij} - \bar{\rho}\widetilde{u_i''u_j''} \right), \quad (2.15)$$

where the term $-\bar{\tau}_{ij}$ represents the compressible Reynolds stress tensor, which arise from the averaging process.

The use of Favre averaging simplifies the governing equations by eliminating additional terms involving density fluctuations, making it the standard approach for compressible turbulence modelling.

Closure Problem

The introduction of the Reynolds stress tensor leads to the closure problem, as it introduces additional unknowns that cannot be directly determined from the averaged equations. To close the system, turbulence models are required.

A common approach is based on the Boussinesq hypothesis, which assumes that the Reynolds stresses can be related to the mean velocity gradients through a turbulent (or eddy) viscosity μ_t . In compressible form, this relation is written as

$$\tau_{ij}^F = -\rho\widetilde{u_i''u_j''} = \mu_t \left(\frac{\partial\tilde{u}_i}{\partial x_j} + \frac{\partial\tilde{u}_j}{\partial x_i} \right) - \frac{2}{3}\mu_t \frac{\partial\tilde{u}_k}{\partial x_k} \delta_{ij} - \frac{2}{3}\bar{\rho}k\delta_{ij}, \quad (2.16)$$

where k is the turbulent kinetic energy and δ_{ij} is the Kronecker delta.

Turbulence Model

In order to determine the turbulent viscosity μ_t , additional transport equations must be solved. In the present work, the k - ω SST turbulence model is employed, which combines the advantages of the k - ω formulation near walls with the robustness of the k - ϵ model in the free stream. It solves two additional transport equations for the turbulent kinetic energy k and the specific dissipation rate ω .

The use of a RANS approach significantly reduces the computational requirements compared to higher-fidelity methods, while still capturing the main effects of turbulence on the mean flow field.

2.2.6 rhoCentralFoam

The numerical simulations are performed using the **rhoCentralFoam** solver available within the OpenFOAM CFD framework. This solver was selected due to its suitability for modelling compressible, high-speed flow regimes where shock waves, expansion fans and strong pressure gradients are present. Furthermore, **rhoCentralFoam** was employed in previous studies carried out within the same research group, which facilitates direct comparison of results.

rhoCentralFoam [19] is a compressible, transient, density-based solver designed for the simulation of high-speed flows involving significant variations in density, pressure and temperature. It is formulated within the finite-volume framework and employs explicit

time integration for the advancement of the governing equations, making the solution subject to a CFL stability constraint.

The solver is based on the central-upwind schemes developed by Kurganov and Tadmor [15] for hyperbolic conservation laws. In contrast to classical upwind methods that rely on the exact or approximate solution of Riemann problems at cell interfaces, central-upwind schemes avoid the explicit construction of Riemann solvers. Instead, they estimate local wave propagation speeds at each cell face and use these to construct numerical fluxes as a combination of left- and right-moving contributions. This approach introduces numerical dissipation in a controlled manner, sufficient to stabilise the solution near discontinuities while maintaining high resolution in smooth regions of the flow. The scheme can be interpreted as a blend between central differencing and upwind methods, where the amount of upwinding is dynamically adjusted based on local wave speeds.

More specifically, the numerical flux at a cell interface is computed as

$$\mathbf{F}_{i+\frac{1}{2}} = \frac{a^+ \mathbf{F}(\mathbf{U}_L) - a^- \mathbf{F}(\mathbf{U}_R)}{a^+ - a^-} + \frac{a^+ a^-}{a^+ + a^-} (\mathbf{U}_R - \mathbf{U}_L), \quad (2.17)$$

where $\mathbf{U}_L$ and $\mathbf{U}_R$ denote the reconstructed conservative variables on the left and right sides of the interface, respectively, and $\mathbf{F}(\mathbf{U})$ is the corresponding physical flux. The quantities a^+ and a^- represent the local maximum and minimum wave propagation speeds, typically defined as

$$a^+ = \max(0, \lambda_{\max}(\mathbf{U}_L), \lambda_{\max}(\mathbf{U}_R)), \quad a^- = \min(0, \lambda_{\min}(\mathbf{U}_L), \lambda_{\min}(\mathbf{U}_R)), \quad (2.18)$$

where $\lambda_{\max}$ and $\lambda_{\min}$ are the largest and smallest eigenvalues of the flux Jacobian (i.e., the characteristic wave speeds of the system). This formulation allows accurate capturing of shock waves and contact discontinuities without excessive numerical diffusion, while maintaining robustness and simplicity compared to traditional Godunov-type methods.

`rhoCentralFoam` supports the inclusion of turbulence models and heat transfer processes, although the latter was not used in the present study.

It should be noted that in more recent versions of OpenFOAM, the solver library has undergone substantial restructuring. The `rhoCentralFoam` application is no longer maintained as a standalone solver in the latest releases. Its functionality has been incorporated into a modular framework for compressible flows, such as the `shockFluid` class introduced in OpenFOAM v11 which provides a more general and extensible structure for density-based and pressure-based algorithms alike. Due to these changes the OpenFOAM versions used for the present study were older such as v6, v2312 and v2412.

2.3 DSMC: Direct simulation Monte Carlo

The kinetic theory of gases provides the fundamental bridge between the microscopic (molecular) description of a gas and its macroscopic, continuum behaviour. Macroscopic properties such as temperature, pressure and density emerge from the collective effects of molecular motion and intermolecular collisions. Gas molecules possess several energy modes: translational, rotational, vibrational and electronic, each characterized by its own excitation energy and relaxation time. Under near-equilibrium conditions, frequent collisions enable efficient exchange of energy between these modes, thereby maintaining thermodynamic equilibrium.

In the continuum formulation, the discrete molecular nature of the gas is neglected and the flow is treated as a continuous medium whose macroscopic properties obey the conservation of mass, momentum and energy. Closure of the governing Navier-Stokes




Figure 2.1: Flow macroscopic properties and microscopic particle modes. Reproduced from [5]

equations requires constitutive relations (e.g. $\tau \sim \nabla \bar{u}$, $\bar{q} \sim \nabla T$) that describe the response of the fluid to gradients in velocity and temperature. These relations rely on the assumption that departures from equilibrium are small and that the flow remains within the near-continuum regime [8].

In rarefied environments, however, such as high-altitude flight or plume expansion into near-vacuum conditions (as in the present study) the molecular mean free path λ becomes comparable to, or larger than, the characteristic flow length scale L . Under these circumstances, intermolecular collisions become infrequent and the different molecular energy modes may no longer equilibrate. As a consequence, strong thermal and chemical nonequilibrium can arise. The continuum assumptions underlying the Navier-Stokes equations, including the validity of constitutive relations, break down and the flow can no longer be treated as a continuum medium.




Figure 2.2: Physical models and flow regimes in terms of Kn . Reproduced from [5]

In rarefied cases, kinetic description is required. A fully microscopic treatment of a gas would require tracking every individual molecule, an impossible task because of the

enormous number of particles involved. Instead, kinetic theory relies on the statistical description provided by the distribution function. The Boltzmann equation governs the evolution of the molecular velocity distribution function $f(t, \bar{x}, \bar{u})$, providing a statistical representation of the gas at the microscopic level:

$$\left(\frac{\partial}{\partial t} + \bar{u} \cdot \nabla_{\bar{x}} + \bar{F} \cdot \nabla_{\bar{u}} \right) f(t, \bar{x}, \bar{u}) = \left(\frac{\partial f}{\partial t} \right)_{\text{coll}}. \quad (2.19)$$

This integro-differential equation describes how the number of molecules within a phase-space element $d\bar{x} d\bar{u}$ changes due to three fundamental mechanisms: (i) convection of particles through the spatial element $d\bar{x}$ as they move; (ii) changes in molecular velocity caused by external forces $\bar{F}$ acting per unit mass; and (iii) depletion and replenishment of particles in the velocity interval $\bar{u} + d\bar{u}$ as a result of intermolecular collisions. The last of these is represented by the nonlinear collision integral $\left(\frac{\partial f}{\partial t} \right)_{\text{coll}}$, which is responsible for driving the gas toward thermodynamic equilibrium. It describes the binary collisions. The collision integral takes the following form:

$$\left(\frac{\partial f}{\partial t} \right)_{\text{coll}} = \int_{-\infty}^{\infty} \int_0^{4\pi} n^2 (f^* f_1^* - f f_1) u_r \sigma_T d\Omega d\mathbf{u}_1 \quad (2.20)$$

Where the superscript $*$ represent post-collision properties, f and f_1 are the velocity distribution function at u and u_1 respectively, u_r is the relative speed of two colliding molecules, σ_T is the molecular cross-section, and Ω is the solid collision angle. Solving the Boltzmann equation is cumbersome due to the complexity of this integral term.

Among the numerical methods derived from the Boltzmann equation, the **Direct Simulation Monte Carlo (DSMC)** technique has become the standard tool for the simulation of rarefied, nonequilibrium gas flows. A formal mathematical proof published in 1992 [27] demonstrated that, in the limit of an infinite number of simulated particles, the DSMC method converges to the solution of the Boltzmann equation for a monatomic gas. This result firmly established the theoretical foundation of DSMC, which is described extensively in Bird's 1994 monograph [4].

DSMC is a stochastic, particle-based method in which each simulation particle represents a large number of real gas molecules, thereby reducing the computational cost compared with fully deterministic approaches such as molecular dynamics (MD).

dsmcFoam

In this work, the DSMC simulations are carried out using the **dsmcFoam** solver, released with OpenFOAM 1.7. Although its purpose is comparable to established DSMC codes such as MONACO, DAC, or SPARTA, **dsmcFoam** is distinct in that its entire functionality is integrated directly into the OpenFOAM framework. This design ensures a high degree of modularity and facilitates coupling with the wide range of solvers available in OpenFOAM. It takes advantage of OpenFOAM's capabilities, such as particle tracking in various meshes and can handle arbitrary 2D/3D geometries and simulate an arbitrary number of gas species. Additionally, it benefits from unlimited parallel processing. The code is open-source and distributed under the same GPL license as its OpenFOAM base, with current releases publicly available through the project repository [29].

dsmcFoam provides several parameters that can be adjusted to maintain numerical stability and to minimise statistical fluctuations during the simulation. The most relevant controls include:

- **Number of particles:** Increasing the number of simulated particles reduces statistical noise. A common guideline is to maintain at least 5-10 particles per computational cell to ensure reliable density sampling and avoid fluctuations caused by insufficient particle representation.
- **Mesh resolution:** The characteristic mesh size should remain smaller than the local mean free path to capture the flow physics accurately and to prevent sharp variations in particle density between adjacent cells.
- **Time step:** Using a sufficiently small time step limits the number of particles introduced or moved per iteration, resulting in smoother temporal evolution and reduced numerical fluctuations.

2.3.1 Algorithmic Procedure

As with any numerical solver, several inputs must be specified before the simulation can be carried out. These include the computational geometry, the boundary conditions and the molecular interaction model.

In the DSMC method, the gas is represented by a collection of computational particles F_n , each acting as a surrogate for a large number of real molecules sharing similar positions and velocities $\{\mathbf{r}_i, \mathbf{v}_i\}$ for $i = 1, \dots, n$. A key assumption underlying the method is that, over sufficiently small time intervals, molecular transport and collisions can be treated as two independent processes: free molecular motion followed by collisional relaxation.

As illustrated in Fig. 2.3, the DSMC algorithm proceeds by advancing all simulation particles through a discretized computational domain, while tracking their trajectories and internal states. Particle motion is computed deterministically based on the molecular velocities and the chosen time step, whereas intermolecular collisions are handled through stochastic models. According to kinetic theory, the probability of collision between two molecules is proportional to the product of their relative velocity and the molecular collision cross-sectional area. Within each computational cell, microscopic particle information is sampled and averaged to obtain macroscopic flow quantities.

Because DSMC resolves gas behaviour through representative particles, it readily accommodates detailed collision models that are extremely difficult, or even impossible, to incorporate directly into the Boltzmann equation. Unlike numerical schemes that attempt to solve Eq. 2.19 using simplified collision operators, DSMC can naturally reproduce temperature-dependent transport properties such as viscosity, thermal conductivity and diffusion coefficients for real gases.

Another significant advantage is that the computational cost of DSMC does not scale with the number N of chemical species in the mixture. This stands in contrast to continuum solvers, which require solving one transport equation per species, leading to an overall cost that typically increases as N^2 [5]. The situation becomes even more demanding for master-equation solvers, which must additionally integrate a conservation equation for every internal molecular state. These characteristics collectively make DSMC a powerful and versatile approach for modelling rarefied and nonequilibrium gas flows.

During each discrete time step Δt the following sequence of operations is carried out:

1. **Particle Movement:** Each particle is moved according to its velocity, using OpenFOAM's inbuilt particle tracking algorithm, which handles motion of particles across faces of the mesh (and also deals with boundaries). In its mathematical form, the move step for the i th particle is given by:




Figure 2.3: DSMC algorithm flowchart. Reproduced from [29]

$$\mathbf{r}_i(t + \Delta t) = \mathbf{r}_i(t) + \mathbf{v}_i(t) \Delta t,$$

with boundary conditions applied immediately when needed. Any particle that reaches a boundary or a surface has its position and velocity reset accordingly.

2. **Cell Sorting:** After transport, the particles are sorted into computational cells, updating the list of particles in each computational cell. The cell size is typically chosen so that its characteristic dimension is not larger than the mean free path, ensuring that the local equilibrium assumptions used in the collision step remain valid.
3. **Collision Selection:** Only particles within the same cell are considered for collisions, all pairs are candidate collision partners regardless of their actual trajectories. Instead of evaluating all possible pairs directly, which would be prohibitively expensive, DSMC relies on probabilistic sampling.
4. **Collision Execution:** Once a candidate pair is selected, post-collision velocities are assigned using the chosen molecular interaction model.
5. **Sampling of Macroscopic Quantities:** Flow properties such as density, velocity, temperature and heat flux are obtained by averaging particle quantities over each cell.
6. **Update time:** The above operations are repeated with new $t = t + \Delta t$ until the simulation end time t_{end} is reached

2.3.2 Collision models

Physical models govern how molecules interact, how momentum and energy are exchanged during collisions, how internal energy modes are populated, how chemical reactions proceed and how particles interact with solid surfaces [4]. In DSMC, these processes are represented through a set of standard modelling approaches, several of which are outlined below:

- **Scattering cross-section:** Determines how molecules deflect during collisions and thus controls the post-collision velocity distribution.
 - **Hard Sphere (HS):** A simple, idealized model in which molecules are treated as rigid spheres with a fixed diameter. As illustrated in Figure 2.4, the repulsive force becomes effectively infinite upon contact, producing an instantaneous deflection.
 - **Variable Hard/Soft Sphere (VHS/VSS):** More realistic models in which the effective molecular diameter varies with relative speed. These models better reproduce transport properties such as diffusion and viscosity and provide more accurate scattering behaviour than the HS model.
- **Collision frequency:** Dictates how often collisions are sampled within each cell.
 - **No-Time-Counter (NTC) scheme:** A widely used algorithm in which the number of attempted collisions is estimated without explicitly counting molecular pairs. It provides efficient and statistically correct sampling.




Figure 2.4: Typical intermolecular force field and the Hard Sphere model in blue. Adapted and modified from [4]

- **Majorant Collision Frequency (MCF):** An improved technique that uses an upper bound on the collision rate to guarantee stability and simplify the selection of collision partners, particularly in mixtures or highly non-uniform flows.
- **Simplified Bernoulli trails (SBT):** addresses the drawback of repeated collisions in the cells and requires only 2-3 particles per cell to obtain more accurate results, significantly reducing the amount of computational resources required.
- **Internal energy exchange:** Determines how translational energy is redistributed into rotational and vibrational modes during collisions.
 - **Larsen-Borgnakke (LB) model:** A stochastic relaxation model that redistributes energy between translational and internal modes according to equilibrium probabilities, controlled by specified relaxation numbers.
 - **Variable rotational energy exchange probability model of Boyd [7]**
 - **Variable vibrational energy exchange probability model of Vijayakumar et al [26]**
- **Chemical reaction models:** Define the probability that a collision leads to dissociation, recombination, or other chemical transformations.
 - **Total Collision Energy (TCE):** The most commonly used model, in which reactions occur if the total post-collision energy exceeds a threshold.
 - **Vibrationally Favoured Dissociation (VFD):** Enhances the likelihood of reaction for molecules with high vibrational energy, reflecting physical behaviour at high temperatures.

- **Quantum-Kinetic (QK):** A more detailed model incorporating quantized internal states and state-dependent reaction probabilities.
- **Gas-surface interactions:** Describe how particles reflect from solid boundaries.
 - **Specular reflection:** The incident particle reflects like a light ray from a mirror, preserving the tangential velocity component.
 - **Diffuse reflection:** The particle re-emits from the surface with a Maxwellian distribution, losing memory of its incident direction.
 - **Maxwell model:** A combination of specular and diffuse reflection, controlled by an accommodation coefficient that accounts for microscopic surface roughness.

2.4 Hybrid CFD-DSMC

Hybrid approaches for multiscale gas flow simulation are commonly grouped into three main categories, Fig. 2.5-2.7, based on the strategy used to exchange information between the microscopic (DSMC) and macroscopic (CFD) descriptions:

- **Domain-decomposition methods (DDM):** DSMC and CFD solvers operate in different spatial regions and are coupled through an interface. This method is the one described in the previous chapters.




Figure 2.5: DDM. Reproduced from [32]

- **Heterogeneous multi-scale methods (HMM):** A macroscopic solver is applied over the whole domain, while a microscopic solver is invoked locally to supply missing physical information.
- **Internal-flow multiscale methods (IMM):** The microscopic model is embedded within the macroscopic solver, often to directly compute fluxes or constitutive relations.

An alternative approach involves the development of unified solvers capable of operating across the entire flow regime, such as the **unified gas-kinetic scheme (UGKS)** and




Figure 2.6: HMM. Reproduced from [32]




Figure 2.7: IMM. Reproduced from [32]

the **discrete unified gas-kinetic scheme** (DUGKS). These methods solve for the gas distribution function using discrete molecular velocity space, thereby providing a consistent description from continuum to rarefied conditions.

A further class of approaches consists of **particle-particle hybrid methods**, which have been shown to be particularly efficient for high-speed gas flow simulations, especially when complex physical and chemical processes must be modelled.

In this thesis, a domain-decomposition framework is employed, allowing DSMC and CFD to be applied in regions corresponding to their respective validity ranges.

Over the past decades, a wide variety of hybrid algorithms have been developed to address kinetic-continuum coupling. These methods differ in their numerical formulation, operating on particles, Partial Differential Equations (PDEs), or directly on velocity distribution functions. Despite these differences, two fundamental challenges are common to most hybrid approaches.

The first challenge arises from the distinct nature of the variables solved by kinetic and continuum models. Kinetic descriptions evolve the particle velocity distribution function, whereas continuum formulations rely on its macroscopic moments. While the computation of moments from a known distribution is straightforward, the inverse operation, commonly referred to as the *moment problem*, is generally ill-posed. To address this issue, hybrid methods often employ reconstruction techniques such as Chapman-Enskog (CE) expansions, Grad-type distributions or Maximum-Entropy-Distribution (MED) .

The second challenge concerns the definition of reliable switching or breakdown criteria, which must ensure both computational efficiency and physical accuracy.

In practice, this challenges translate to determining the location of the interface, the selection of an appropriate coupling strategy, and the methodology used for information exchange between the DSMC and CFD domains. These aspects are discussed further in Sec. 2.4.1 and Sec. 2.4.2.

2.4.1 Coupling Strategies and Information Exchange in DDM

The choice of coupling strategy is strongly influenced by the flow regime under consideration, particularly the characteristic flow speed and the disparity of timescales between molecular and continuum descriptions. [32]

Explicit Coupling for High-Speed Flows: In high-speed compressible flows, such as hypersonic external aerodynamics or atmospheric entry problems, the characteristic physical timescale is comparable to both the compressible CFL time step of the CFD solver and the mean collision time governing the DSMC method. Under these conditions, both solvers can be advanced synchronously in physical time, making time-accurate (*explicit*) coupling an appropriate choice. The DSMC and CFD solvers are advanced sequentially at each global time step, with interface data exchanged after every update. The information transfer may be based on either flux quantities (mass, momentum, and energy fluxes) or state variables (density, velocity, and temperature). For high-speed flows, flux-based coupling is feasible because the magnitude of the physical fluxes is large compared to the statistical fluctuations inherent in DSMC sampling, resulting in acceptable signal-to-noise ratios.

Implicit Coupling for Low-Speed Flows: In low speed flows, such as those encountered in micro-scale gas dynamics and MEMS, incompressible or weakly compressible CFD formulations are typically employed, allowing relatively large time steps in the continuum solver. However, the DSMC time step remains constrained by the molecular collision time, which is orders of magnitude smaller. Advancing the DSMC solver explicitly to match the macroscopic flow timescale therefore becomes computationally impractical. Moreover, flux-based coupling becomes unreliable because the physical fluxes across the interface are small, often of the same order as DSMC statistical noise. Instead, *implicit* coupling approaches are preferred, where steady-state solutions of the molecular and continuum subdomains are obtained iteratively. A commonly adopted procedure is the **Schwarz alternating method**, in which boundary conditions are exchanged repeatedly between the two solvers over an overlap region until convergence is achieved. state variables (Dirichlet conditions) or gradients (Neumann conditions) are typically transferred across the interface, as these quantities exhibit lower statistical uncertainty than instantaneous fluxes. The iterative nature of the method effectively decouples the molecular and continuum timescales while still ensuring a consistent solution.

State-Based and Flux-Based Interface Coupling

As mentioned before, two primary approaches are commonly used to exchange information across the continuum-molecular interface: flux-based coupling and state-based coupling, represented in Fig.2.8. The distinction between these approaches lies in the type of physical quantities enforced at the interface.

Flux-based Coupling: Conservation fluxes are directly transferred between the molecular and continuum solvers. At the interface, DSMC particles crossing the boundary provide instantaneous estimates of mass, momentum, and energy fluxes. These fluxes are computed by tracking particle crossings and accumulating their contributions over a sampling interval. The resulting averaged fluxes are then imposed as boundary fluxes on the adjacent Navier-Stokes cells. In the opposite direction, the continuum solver evaluates the same flux quantities from macroscopic gradients of flow variables, such as velocity and temperature. These continuum-derived fluxes are subsequently used to construct particle distributions in the DSMC interface cells. Although this method enforces conservation in a direct and physically intuitive manner, it is sensitive to the inherent statistical fluctuations




Figure 2.8: Schematic of typical hybrid molecular-fluid coupling methods. Reproduced from [10]

of DSMC. The instantaneous particle flux at the interface can exhibit significant scatter. As a consequence, flux-based coupling typically requires extensive temporal averaging and a large number of simulation particles to reduce statistical noise to acceptable levels.

State-based Coupling: Also known as **Dirichlet-Dirichlet coupling**, exchanges macroscopic state variables across the interface. On the molecular side, particle properties within DSMC cells adjacent to the interface are averaged to obtain density, velocity, and temperature. These averaged quantities are then imposed as Dirichlet boundary conditions for the Navier-Stokes solver. Conversely, the continuum solution provides the same set of macroscopic state variables at the interface, which are used to generate particles in the DSMC boundary cells. Typically, particles are sampled from velocity distributions consistent with the prescribed density, bulk velocity, and temperature. In this way, continuity of thermodynamic state is enforced, while fluxes emerge naturally from the solution rather than being prescribed directly. This method often achieves stable coupling with fewer DSMC samples.

Several numerical techniques have been developed to treat the coupling at the kinetic-continuum interface. Four commonly used approaches can be identified.

The first approach is the Marshak condition, a flux-based method originally introduced in radiative heat transfer. In this technique, the incoming and outgoing particle half-fluxes¹ at the interface are matched to ensure conservation of mass, momentum, and energy. The net flux is obtained by combining these contributions. Assuming a Maxwellian distribution² on the continuum side, the macroscopic interface properties can be determined.

The second approach consists of extrapolating macroscopic flow properties, such as density, velocity, and temperature, from the continuum region toward the interface. This method belongs to the class of state-based coupling strategies, and its accuracy depends on the order of the extrapolation scheme.

The third approach involves the extrapolation of net fluxes instead of state variables.

¹Half-fluxes refer to flux contributions associated with particles whose velocity components are directed either toward or away from the interface, typically considering only half of the velocity space normal to the boundary.

²A Maxwellian distribution describes a gas in local thermodynamic equilibrium, where particle velocities follow a Gaussian distribution determined by the local macroscopic density, velocity, and temperature.

While similar in spirit to property extrapolation, this method focuses directly on conserved quantities and is therefore more closely related to flux-based coupling.

The fourth approach is based on the asymptotic solution of the linear kinetic half-space problem³. This treatment is particularly useful in situations where strong non-equilibrium effects occur and simpler equilibrium-based assumptions are no longer valid.




Figure 2.9: Flowchart of the coupled NS-DSMC method. Adapted from [2]

Considerations for axi-axisymmetry: It is necessary to set two additional parameters due to the onset of the axisymmetric simulation: radial extension domain and maximum radial weighting factor (MaxRWF). The former indicates that the domain size is in the Y-direction and the latter is used to adjust the real number of molecules represented by the DSMC particles in the near and far axis. This is because there are usually many particles in the cells distant from the axis but very few next to the axis in an axisymmetric simulation.

³The linear kinetic half-space problem describes the behavior of a dilute gas near an interface by solving a simplified kinetic equation in a semi-infinite domain, often used to derive accurate boundary conditions under non-equilibrium conditions.

2.4.2 Breakdown parameters

Knudsen number

Several parameters are used to characterise the degree of rarefaction and nonequilibrium in a gas. Among them, the most widely used is the **Knudsen number** Kn , defined as the ratio of the molecular mean free path λ to the characteristic length scale of the system L :

$$Kn = \frac{\lambda}{L}$$

It is used as a dimensionless parameter to evaluate whether molecular-level interactions affect continuum-level transport of momentum and energy.

In the limit of vanishing Knudsen number, intermolecular collisions occur so frequently that the gas remains in perfect thermodynamic equilibrium. As the Knudsen number increases, however, collisions become progressively less frequent, and the flow departs from equilibrium. In the extreme free-molecular limit, intermolecular collisions are virtually absent and gas-surface interactions dominate the behaviour of the system. Because relaxation toward equilibrium is driven by intermolecular collisions, nonequilibrium effects grow stronger with increasing Knudsen number. The main rarefaction regimes are illustrated in Fig. 2.10. A brief summary of the flow regimes and the numerical methods typically




Figure 2.10: Knudsen number regimes, adapted from [11]

employed in each is as follows [11]:

- $Kn \rightarrow 0$: Inviscid (continuum) regime, governed by the Euler equations.
- $Kn \leq 0.001$: Continuum regime, Navier-Stokes-Fourier (NSF) equations are valid.
- $0.001 \leq Kn \leq 0.1$: Slip regime, NSF equations remain applicable with velocity-slip and temperature-jump boundary conditions.
- $0.1 \leq Kn \leq 10$: Transition regime, continuum assumptions break down. Kinetic descriptions such as the Boltzmann equation or particle methods (e.g. DSMC) are required.
- $Kn \geq 10$: Free-molecular regime, the collision-less Boltzmann equation or particle methods like DSMC are appropriate.

Bird parameter

Bird [3] proposed a semi-empirical parameter to characterize continuum breakdown in steady, expanding flows. The **Bird parameter** P is defined as

$$P = \frac{U}{\rho\nu} \left| \frac{d\rho}{ds} \right| = M \sqrt{\frac{\pi\gamma}{8}} \frac{\lambda}{\rho} \left| \frac{d\rho}{ds} \right| \quad (2.21)$$

where U is the local flow velocity, ν is the molecular collision frequency, M is the local Mach number, γ is the ratio of specific heats, and s denotes the distance measured along a streamline. This parameter relates the local flow acceleration to the rate of intermolecular collisions and provides an estimate of the degree of nonequilibrium in the flow.

Generalised Bird parameter

Although the study presented in [3] suggests that a value of $P \approx 0.05$ provides a reasonable criterion for continuum breakdown in steady expanding flows, this threshold is not always sufficient for more complex flow configurations [28]. In such cases, density alone may not adequately describe the onset of nonequilibrium effects.

Following the same reasoning, the breakdown parameter can be generalized by replacing density with any macroscopic flow quantity Q , such as temperature or velocity. This allows the formulation of alternative criteria that account for gradients in different flow properties and may provide a more reliable indication of continuum breakdown in multidimensional and highly non-uniform flows.

$$P_Q = M \sqrt{\frac{\pi\gamma}{8}} \frac{\lambda}{Q} \left| \frac{dQ}{ds} \right| \quad (2.22)$$

Since the evaluation of the stream-wise gradient required to compute the breakdown parameter P_Q depends on the local velocity components, its definition becomes problematic in regions where the velocity approaches zero, such as at stagnation points.

Gradient-length local Knudsen number

An extensive numerical investigation of one-dimensional normal shock waves and two-dimensional bow shocks conducted by Boyd et al. [6] concluded that the **gradient-length local Knudsen number** $\mathbf{Kn}_{GLL}$ provides a better indication of continuum breakdown than P for hypersonic compressed flows. Where L instead of an overall flow dimension is the gradient length scale of a macroscopic quantity Q , where Q can represent the local velocity, density, pressure, temperature, viscosity etc. $L = \frac{Q}{|\nabla Q|}$ *e.g.* : $L = \frac{\rho}{d\rho/dx}$

$$Kn_{GLL} = \frac{\lambda}{Q} \left| \frac{dQ}{dx} \right| \quad (2.23)$$

A large Knudsen number occurs with either a large mean free path or a small scale length of the macroscopic gradients. The former is the case in this study and it is due to a very low gas density. While the alternative requirement of small characteristic dimension can be met at any density. When calculating based on velocity magnitude, the gradient is normalized by the maximum of the local velocity magnitude and the local speed of sound. It is generally assumed that continuum breakdown occurs whenever $\mathbf{Kn}_{GLL}$ is greater than 0.05. Eqs 2.22 and 2.23 are not independent and have the relation:

$$\frac{P_Q}{Kn_Q} = M \sqrt{\frac{\pi\gamma}{8}} \cos \theta, \quad (2.24)$$

where θ is the angle between the gradient of the macroscopic quantity ∇Q and the local flow direction. Since $\sqrt{\frac{\pi\gamma}{8}}$ is smaller than unity for most gases, the value of P_Q is always lower than Kn_Q in subsonic flow regions.

In the near-wall region, both the Mach number M and the directional cosine $\cos \theta$ decrease toward zero. As a consequence, P_Q becomes several orders of magnitude smaller

than Kn_Q in the immediate vicinity of solid surfaces. This behaviour limits the usefulness of P_Q as a breakdown indicator near stagnation points and walls, where velocity magnitudes and streamwise gradients are small.

Maximum gradient-length local Knudsen number & Maximum Bird's parameter

In 2003, Wang and Boyd [28] investigated the breakdown of the Navier-Stokes equations in hypersonic viscous flows over several geometries, such as a sharp cone and a hollow cylinder-flare configuration. The objective of their study was to identify the conditions under which the continuum assumption becomes invalid for hypersonic gas flows.

Their analysis was based on a detailed comparison between numerical results obtained using CFD methods and DSMC approach. As part of this work, Wang and Boyd introduced a new parameter, denoted as maximum gradient-length local (GLL) Knudsen number $Kn_{\max}$, which represent a modified form of Bird's original Kn_{GLL} criterion and also compared it to $P_{\max}$ from Bird's parameter.

$$P_{\max} = \max(P_\rho, P_T, P_U) \quad (2.25)$$

$$Kn_{\max} = \max(Kn_\rho, Kn_T, Kn_U) \quad (2.26)$$

They showed that when $Kn_{\max}$ exceeds a value of approximately 0.05, the difference between full DSMC and full CFD exceeds 5 %, indicating the onset of continuum breakdown.

Parameter B

Similarly, Garcia et al. [12] use a continuum breakdown parameter derived directly from the higher-order terms of the **Chapman-Enskog** expansion of the distribution function. This parameter, denoted by **B** (Eq. 2.27), incorporates normalized shear stress (Eq. 2.28) and heat flux contributions (Eq 2.29):

$$B = \max(|\tau_{i,j}^*|, |q_i^*|) \quad (2.27)$$

$$\tau_{i,j}^* = \frac{\mu}{p} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} - \frac{2\partial u_k}{3\partial x_k} \delta_{i,j} \right) \quad (2.28)$$

$$q_i^* = -\frac{\kappa}{p} \sqrt{\frac{2m}{\kappa T}} \frac{\partial T}{\partial x_i} \quad (2.29)$$

where indices i, j, and k stand for the three velocity components. A conservative switching criterion of $B > 0.1$ has been suggested, whereby DSMC is employed when this threshold is exceeded and CFD is used otherwise.

Tiwari's criterion

Tiwari [23] proposed an alternative switching criterion based on the magnitude of a non-equilibrium parameter, denoted by $\|\phi\|$, defined as

$$\|\phi\| = \frac{1}{\rho RT} \left[\frac{2}{5} \frac{|q|^2}{RT} + \frac{1}{2} \|\tau\|^2 \right]^{1/2} \quad (2.30)$$

where q is the heat flux vector, $\|\tau\|$ is the Euclidean norm of the stress tensor τ , ρ is the local mass density, R is the universal gas constant, and T is the translational temperature. For the continuum approximation to remain valid, this parameter must be much smaller than unity. Meaning that kinetic solvers should be employed when $\|\phi\| > \epsilon$, with $\epsilon \ll 1$.

R13 based criteria

Lockerby et al. [17] proposed an alternative breakdown criterion based on the discrepancy between the constitutive relations used in the Navier-Stokes equations (Newtonian stress and Fourier heat flux) and those obtained from the higher-order Regularized 13-Moment (R13) equations. The corresponding breakdown parameters are defined as

$$\text{Kn}_\tau = \frac{|\tau_{ij}^{(\text{R13})} - \tau_{ij}^{(\text{NS})}|}{|\tau_{ij}^{(\text{NS})}|}, \quad \text{Kn}_q = \frac{|q_i^{(\text{R13})} - q_i^{(\text{NS})}|}{|q_i^{(\text{NS})}|}. \quad (2.31)$$

The R13 equations belong to a class of extended hydrodynamic models derived from kinetic theory, in which additional transport equations are introduced for higher-order moments of the molecular velocity distribution. While the Navier-Stokes equations assume that stress and heat flux are linearly related to velocity and temperature gradients (i.e., Newton's law of viscosity and Fourier's law of heat conduction), the R13 formulation accounts for non-equilibrium effects by including higher-order corrections to these relations.

As a result, the R13 model provides a more accurate description of gas behaviour in transitional and rarefied regimes, where the assumptions underlying the Navier-Stokes equations begin to break down. The parameters Kn_τ and Kn_q therefore quantify the relative error introduced by using Navier-Stokes constitutive relations instead of a higher-order kinetic-based model.

An important advantage of this approach is that estimates of the R13 stress and heat flux terms can be constructed using only Navier-Stokes-level information, allowing the breakdown parameter to be evaluated without solving the full R13 system.

In the literature, values of Kn_τ and Kn_q on the order of 0.1 are typically considered indicative of significant non-equilibrium effects, while lower thresholds (e.g., 0.01-0.05) may be used as more conservative criteria for the onset of breakdown.

Tsien's number

In supersonic and hypersonic flows rarefaction effects are often characterized using the parameter introduced by Tsien [24]. This parameter was derived in the context of compressible boundary layers in slip flow conditions near solid surfaces.

Tsien considered the ratio between the molecular mean free path l and the boundary-layer thickness δ ,

$$\frac{l}{\delta} = \frac{l}{L} \frac{L}{\delta}, \quad (2.32)$$

where L is a characteristic body dimension. Using relations from kinetic theory, the ratio l/L can be approximated as $\frac{l}{L} \sim \frac{\nu}{aL}$, where ν is the kinematic viscosity and a is the speed of sound. The definitions of Mach and Reynolds numbers, $M = \frac{U}{a}$ and $Re = \frac{UL}{\nu}$, leads to the scaling

$$\frac{l}{\delta} \sim \frac{L}{\delta} \frac{\nu}{aL} \sim \frac{L}{\delta} \frac{M}{Re}. \quad (2.33)$$

The resulting expression depends on the ratio L/δ , which is determined by the boundary-layer scaling. For high Reynolds numbers, where $\delta/L \sim Re^{-1/2}$, the parameter becomes

$$\frac{l}{\delta} \sim \frac{M}{\sqrt{Re}}, \quad (2.34)$$

while in the opposite limit of small Reynolds numbers, $\delta/L \sim 1$, the scaling reduces to

$$\frac{l}{\delta} \sim \frac{M}{Re}. \quad (2.35)$$

Tsien suggested that slip-flow effects become significant when

$$0.01 < \frac{l}{\delta} < 1. \quad (2.36)$$

Because this parameter is based on the boundary-layer thickness δ , its practical application is mainly limited to flows where a well-defined boundary layer exists. For strongly expanding flows, such as nozzle plumes in low-pressure environments, the absence of a classical boundary layer makes the direct use of Tsien's parameter less straightforward. Nevertheless, Tsien also considered the transition to the free-molecule regime by examining the ratio l/L . In the original work it is stated that

“If one takes the limiting ratio of l/L to be 10 for the free molecule flow, then this realm of fluid mechanics occupies a region of the $M-Re$ plane given by $l/L \sim M/Re > 10$.”

This relation provides a qualitative criterion for identifying the free-molecule regime in terms of Mach and Reynolds numbers, without reference to the boundary-layer thickness δ . It can be interpreted as an alternative scaling of the Knudsen number, where ($M/Re > 10$) corresponds approximately to $Kn \gg 1$. Since this relation essentially represents a rescaled Knudsen number, its numerical results are not reported here.

Deborah number

Deborah number is the ratio of some molecular relaxation time τ_{relax} and the flow characteristic time τ_{flow} , it originates from rheology but is also used in rarefied gas dynamics.

$$De = \frac{\tau_{relax}}{\tau_{flow}}, \quad (2.37)$$

τ_{relax} can represent different molecular time scalars such as the mean collision time or the rotational, vibrational and electronic relaxation times. Values approaching or exceeding unity indicate strong non-equilibrium behaviour, for which continuum-based methods are no longer expected to be valid.

Wu's thermal non-equilibrium indicator P_{Tne}

Wu et al. [31] introduced a thermal equilibrium breakdown parameter defined, for diatomic gases at moderate temperatures, as the ratio of the difference between translational and rotational temperatures to the translational temperature. This concept can be extended to atomic gases by considering the difference between translational temperatures associated with different molecular motion components. For high-temperature flows, additional internal modes, such as vibrational energy, may also be incorporated into the definition of a thermal non-equilibrium indicator.

In their study, the thermal non-equilibrium parameter is defined as

$$P_{Tne} = \left| \frac{T_{Tr} - T_R}{T_{Tr}} \right|, \quad (2.38)$$

where T_{Tr} and T_R denote the translational and rotational temperatures, respectively. When the local value of P_{Tne} exceeds a prescribed threshold P_{Tne}^{Thr} (chosen as 0.1 in the study), the assumption of thermal equilibrium is considered invalid, indicating that the flow in that region cannot be accurately described using NS solvers.

It should be noted that P_{Tne} can be evaluated only within the DSMC solver, since the CFD solver employed in their work (just like `rhoCentralFoam` used in the present study) does not incorporate multi-temperature modelling. Consequently, this parameter is used in conjunction with the local Knudsen number (or other breakdown indicators) to guide potential switching between DSMC and NS descriptions. In particular, it helps identify regions where thermal non-equilibrium effects diminish during the iterative coupling process, allowing a transition from DSMC back to NS modeling. Also, it may be justified to either use a multiple-temperature NS code, or implement the multiple-temperature capability into HYB3D/`rhoCentralFoam` to promote the computational efficiency by simulating thermal non-equilibrium region without using DSMC.

Levermore matrix

Levermore et al. [16] presented a matrix which can be used to assess the validity of the Navier-Stokes approximation for rarefied gas flows. The matrix is derived from moment realizability conditions for the underlying velocity distribution function and is the simplest such matrix which includes terms related to both stress and heat flux. This matrix, composed of terms intrinsic to the Navier-Stokes equations, should be non-negative definite when continuum fluid dynamics are valid.

$$V_{NS} = I - \frac{\mu}{\rho\theta} [\nabla_x u + (\nabla_x u)^T - \frac{2}{3} \nabla_x \cdot u I] - \frac{2}{3} \frac{\kappa^2}{\rho^2 \theta^3} \nabla_x \theta (\nabla_x \theta)^T \quad (2.39)$$

where I is the identity matrix, μ is the dynamic viscosity, $\theta = k_B T/m$ is proportional to the specific internal energy, κ is the thermal conductivity. The first term in brackets corresponds to the normalized viscous stress contribution, while the second term represents the normalized heat flux contribution. The eigenvalues of V_{NS} provide a measure of continuum breakdown: values significantly deviating from 1 indicate that the Navier-Stokes assumptions may no longer hold.

This matrix is portable and may be applied to any Navier-Stokes calculation where breakdown is suspected. The inclusion of both stress and heat flux effects is an important advantage. The breakdown of the Navier-Stokes approximation of these fluxes is not necessarily correlated, so that having one criterion which incorporates both effects is of considerable interest.

Entropy generation rate

An entropy generation rate has been proposed by Chen et al. [9] as a continuum breakdown parameter since it has a stronger physical foundation than Knudsen numbers.

$$\dot{S}_{gen} = \frac{\tau_\mu}{T} : \nabla_x \mathbf{u} - \frac{\mathbf{q}}{T^2} \cdot \nabla_x T \quad (2.40)$$

where τ_μ is the viscous stress tensor, $\mathbf{q}$ is the heat flux vector, $\nabla_x \mathbf{u}$ and $\nabla_x T$ are the velocity and temperature gradients, respectively, $:$ denotes the double contraction of two tensors. The first term represents entropy production due to viscous dissipation, while the second term corresponds to entropy production due to heat conduction. Large values of $\dot{S}_{gen}$ indicate significant nonequilibrium effects, signalling potential breakdown of continuum assumptions.

The study conducted was hypersonic flow over a hollow cylinder/flare geometry. However, all of the breakdown parameters considered in the paper failed to predict non-continuum behaviour in the shock front.

Chapter 3

System model

The physical problem considered in this study consists of a supersonic jet issuing from a rocket nozzle and expanding into a near-vacuum environment, representative of lunar conditions. Due to the low ambient pressure, the flow undergoes rapid expansion, leading to strong gradients and the formation of shock structures. Furthermore, as the gas expands, regions of the flow may no longer satisfy the continuum assumption, making a purely continuum-based approach insufficient.

CFD simulations are used to model the entire flow and DSMC is employed to model the flow from the nozzle exit. The objective of this work is to evaluate different breakdown criteria for identifying the interface between continuum and rarefied regions by comparing CFD predictions with DSMC results.

The nozzle geometry considered in this study is a conical convergent-divergent (CD) configuration. A short straight extension is included downstream of the divergent section to represent the physical continuation of the nozzle exit used in the experimental setup. The nozzle is oriented perpendicular to the impingement surface.

Three simulation configurations were analysed, all sharing the same nozzle geometry but differing in the stand-off height, defined as the distance between the nozzle exit and the surface. The values considered are approximately 0.04 m, 0.08 m and 0.16 m, as shown in Fig. 3.1 allowing the investigation of plume-surface interaction across different expansion regimes.

The flow is modelled as compressible and transient, with nitrogen treated as an ideal gas. Turbulence effects are accounted for using a Reynolds-averaged Navier-Stokes (RANS) approach, as described in Chapter 2. The simulations aim to capture the main features of the plume expansion, including shock formation, expansion fans and interaction with the surface, while providing data for comparison with DSMC results.

3.1 Numerical Setup

The numerical discretisation and solution strategies were selected to ensure stability and accuracy for compressible, high-speed flow simulations characterised by strong gradients and shock waves.

Temporal discretisation was performed using a first-order explicit Euler scheme, consistent with the `rhoCentralFoam` solver. A variable time-step approach was adopted, in which the time step Δt is dynamically adjusted to maintain the Courant number below a prescribed limit, set to 0.2 for all simulations. The time-step value specified in the `system/controlDict` dictionary is therefore used only as an initial estimate at the start




Figure 3.1: Comparison of the different stand-off heights

of the simulation¹. From the first iteration onward, the solver dynamically adapts Δt based on the local flow velocities and mesh resolution to satisfy the CFL condition. As a consequence, finer meshes require a smaller initial time-step value.

Convective fluxes were computed using the Kurganov central-upwind scheme, which provides robust shock-capturing capabilities without the need for a Riemann solver. Reconstruction of flow variables at cell faces was performed using a van Albada limiter, which offers a good compromise between accuracy and suppression of non-physical oscillations near discontinuities.

Diffusive terms and gradients were discretised using second-order schemes. In particular, a cell-limited linear formulation was adopted for gradient evaluation, which constrains excessive gradient magnitudes and improves solution stability in regions with strong flow gradients.

For the turbulence transport equations, a first-order upwind discretisation was employed for the convective terms in order to maintain boundedness of the solution and prevent non-physical oscillations.

The resulting system of equations was solved using iterative linear solvers. A diagonal solver was used for the density-based variables, while a `smoothSolver` with Gauss–Seidel smoothing was employed for velocity, energy and turbulence quantities. Tight convergence tolerances were specified to ensure accurate resolution of the flow field, particularly in regions affected by strong compressibility effects.

Overall, the numerical setup represents a balance between accuracy and robustness,

¹If the initial Δt is chosen too large, the Courant number constraint may be violated during the first iteration, causing the simulation to terminate immediately, often without explicit error messages.

with the lowest stand of height simulation being most stable and accurate.

3.2 Boundary Conditions

The inlet boundary condition was defined in terms of total pressure and total temperature. To ensure numerical stability, the total pressure was gradually increased² from an initial low value to the target value of 10^6 Pa during the early stages of the simulation.

Far-field boundaries were treated using non-reflecting conditions in order to minimise artificial wave reflections. The impingement surface was modelled as a no-slip wall with a fixed temperature of 300 K. Symmetry conditions were applied in the span-wise direction using a wedge formulation.

The thermal boundary conditions were defined to reproduce, as closely as possible, the experimental configuration.

In the experiments, the characteristic flow timescale was extremely short due to the high-speed nature of the jet. Under such conditions, heat transfer between the gas and the solid walls is negligible, and the wall temperature can reasonably be assumed constant. Consequently, an isothermal wall condition was imposed in the simulations, with a fixed temperature of $T_{wall} = 300$ K.

This assumption reflects the experimental reality in which the thermal inertia of the wall prevents any significant temperature variation during the flow event.

Alternative boundary condition formulations were considered. A zero-gradient temperature condition at the wall could allow thermal adjustment of the fluid adjacent to the surface. However, such a treatment would imply an adiabatic wall, inconsistent with the known ambient laboratory conditions. Similarly, implementing a fully coupled conjugate heat transfer model would require detailed knowledge of material properties and transient thermal behaviour, which were not available.

3.3 Turbulence Modelling

Turbulence effects were modelled using a Reynolds-Averaged Navier-Stokes (RANS) approach. In particular, the Shear Stress Transport (SST) k - ω model was selected due to its robustness in flows with strong adverse pressure gradients, shock interactions, and boundary layer separation, which are characteristic of compressible jet expansion.

The SST formulation combines the advantages of the k - ω model in the near-wall region with the k - ϵ behaviour in the free stream, providing improved accuracy across different flow regimes.

3.3.1 Inlet Turbulence Conditions

Due to the absence of experimental turbulence measurements, inlet turbulence quantities were prescribed using engineering estimates. The turbulent kinetic energy k was defined based on a turbulence intensity of 1%, while the turbulence length scale was estimated

²The rate at which the total pressure was increased was found to have a significant impact on the stability of the simulation. If the increase was too rapid, the simulation became unstable and diverged. Conversely, if the increase was sufficiently gradual to maintain stability, but still quite fast, it led to a non-physical pressure build-up immediately downstream of the inlet, resulting in backflow that persisted throughout the simulation.

from the inlet hydraulic diameter using the empirical relation

$$\ell = 0.07 D_h,$$

where $D_h \approx 0.0084$ m.

Within the SST framework, the specific dissipation rate ω is related to k and the turbulence length scale through

$$\omega = \frac{\sqrt{k}}{C_\mu^{1/4} \ell},$$

with $C_\mu = 0.09$.

Since the inflow velocity is initially low and the flow is primarily driven by pressure gradients, the imposed turbulence levels at the inlet are small. Turbulence is therefore predominantly generated within the nozzle due to shear, compressibility effects and boundary layer development.

3.3.2 Wall Treatment

Solid walls were treated using a low-Reynolds-number formulation of the SST model. In this approach, the viscous sublayer is directly resolved with a sufficiently fine mesh near the wall ($y^+ < 1$).

The turbulent kinetic energy k approaches zero at the wall due to the no-slip condition and a small non-zero value was imposed to ensure numerical stability. The specific dissipation rate ω increases significantly near the wall, consistent with its asymptotic behaviour

$$\omega \sim \frac{\nu}{\beta_1 y^2}.$$

This formulation allows accurate resolution of near-wall turbulence without relying on empirical wall functions.

3.3.3 Outlet Treatment

At outlet boundaries, turbulence quantities were prescribed using flow-direction-dependent conditions. Under normal operation, where flow exits the domain, a zero-gradient condition is effectively applied, allowing turbulent structures to convect out without artificial reflection.

In the presence of local backflow, turbulence quantities are assigned based on internal field values to prevent the introduction of non-physical conditions.

3.4 Mesh Design

This section presents the mesh design used in the first major simulation. Only the description of this case is included since the same procedure was applied when generating the subsequent meshes with different stand-off heights.

One of the key elements contributing to choices made in both geometry and mesh designs is the fact that the **axial-symmetry** of the nozzle can be exploited to reduce computational cost. Instead of modelling the full 3D domain, a narrow wedge is used, with the Y-direction as the radial direction, the X-direction being axial, and the Z-direction being circumferential.

In order to apply the wedge functionality in OpenFOAM the mesh has to have a depth of only one cell in the azimuthal direction, the wedge must be aligned symmetrically

around the central plane of the geometry, and all boundaries in the azimuthal direction should be set as 'wedge' type in the boundary conditions. The mesh must also maintain consistent topology along the wedge faces to avoid non-orthogonal or skewed cells, which can compromise solver stability. All of these requirements have been met in the presented cases.

3.4.1 Domain Geometry




Figure 3.2: Geometry of the initial simulation (full CFD)

The geometric model (Figure 3.2) of the nozzle and the surrounding computational domain was provided in SolidWorks by the team responsible for the experimental setup, ensuring full consistency between the numerical and experimental configurations.

For the wedge functionality mentioned above, a narrow wedge of 5° , symmetric about the central axial plane, is used. This wedge preserves the essential physics of the flow while significantly reducing the number of cells, and its periodic faces will later be defined appropriately in OpenFOAM.

The nozzle itself is a simple conical convergent-divergent (CD) geometry. Surrounding the nozzle, the domain consists of a far-field boundary at the top, an additional far-field outlet to the left of the nozzle, and a wall representing the lunar surface (ground in case of the experiments) on the right-hand side. The lower boundary of the domain reduces to a single line due to the axisymmetric wedge construction. Since the two wedge faces are treated as periodic in OpenFOAM, this boundary is effectively removed in the simulations.

3.4.2 Mesh generation with ICEM CFD

The mesh used in this work was generated in ANSYS ICEM CFD, which is a block-structured meshing tool based on hexahedral topology. ICEM operates by defining one or more blocks that are subsequently manipulated, subdivided, and associated with the geometry until their shape conforms to the domain to be meshed.




Figure 3.3: Full CFD simulation Mesh in ICEM

The meshing process began with a single large hexahedral block that encompassed the entire computational domain. This initial block was then subdivided into smaller blocks to match the main geometric features of the nozzle and the surrounding flow region. Each subdivision was performed so that the block topology aligned with important surfaces and flow regions, such as the convergent-divergent section and the external extension of the nozzle.

Because the geometry included a wedge section (arising from the axial-symmetry treatment), that reduced the bottom of the domain to a single line, the corresponding block vertices at the bottom were merged, effectively collapsing one block side into a line.

Once the block structure matched the geometry, edges of the blocks were associated with the curves of the CAD model, and vertices were projected onto the corresponding points of the geometry. This step forces the block topology to adopt the exact shape of the domain and guarantees that the final hexahedral elements follow the nozzle contour and the outer boundaries accurately.

To resolve the steep gradients expected inside the nozzle, a boundary-layer-type refinement was generated using an O-grid, a standard ICEM technique in which a circular layer of blocks is wrapped around a surface. The O-grid produces a smooth clustering of cells normal to the wall, improving resolution of near-wall features and accelerating convergence. The initial first layer height and the boundary layer thickness estimations are presented in the next subsection [3.4.3](#).

For the lunar surface boundary layer, where no natural O-grid could be wrapped, mesh refinement was achieved by manipulating node distribution along the block edges. This is thanks to the fact that ICEM allows control over node spacing and growth ratios, enabling the placement of a fine, smoothly stretched mesh near the surface without modifying the overall block topology.

The final mesh obtained from this procedure in the case of the shortest stand off height is shown in [Figure 3.3](#).

In Figure 3.4 the depth of one cell can be observed for the periodic wedge function.




Figure 3.4: Isometric view of the full CFD simulation Mesh in ICEM

The final mesh of this geometry has 28838 cells in total, 28489 of which are hexahedra and 358 are prisms, with the total Volume of $3.44869 \times 10^{-6} m^3$.

3.4.3 Estimation of the First Boundary-Layer Height

In the present study, the computational mesh is generated such that the non-dimensional wall distance, y^+ , of the first cell adjacent to the wall is close to one. This resolution ensures that the viscous sublayer of the turbulent boundary layer is fully resolved, allowing direct integration of the turbulence model equations down to the wall without the use of wall functions [22]. Adopting $y^+ \approx 1$ provides a more accurate representation of near-wall velocity and temperature gradients, which is particularly important in regions with strong pressure gradients or flow separation. Although this approach increases the overall number of cells and consequently the computational cost, it improves the physical fidelity of the results and eliminates additional modelling uncertainties associated with wall-function approximations.

The flow conditions and gas properties are summarised in Table 3.1.

An approximate estimation of the first-layer height, y_1 , is carried out to provide an initial guideline for the mesh design. The value will be subsequently refined, if required, based on the y^+ distribution obtained from the initial CFD results. It is important to note that the static temperature at the nozzle exit, $T_e = 300$ K, is assumed to correspond to the flow inside the nozzle (which is considered equal to the temperature outside) prior to complete expansion into the low-pressure external region. The external pressure is $p_{\text{ext}} = 1.8$ Pa, while the internal exit pressure generally remains higher due to imperfect expansion.

Table 3.1: Given flow and thermo-physical parameters.

Quantity	Symbol	Value
Molecular weight	M	28.0134 g/mol
Dynamic viscosity	μ	1.79×10^{-5} kg/(m · s)
Prandtl number	Pr	0.718
Total pressure (inlet)	p^0	1.0×10^6 Pa
External pressure	p_{ext}	1.8 Pa
Total temperature (inlet)	T^0	780 K
Static temperature (exit, internal)	T_e	300 K
Gas constant	R	298 J/(kg · K)
Specific heat ratio	γ	1.4

Step 1: Exit Mach Number

Assuming isentropic expansion between the inlet stagnation and the internal exit plane,

$$\frac{T^0}{T_e} = 1 + \frac{\gamma - 1}{2} M_e^2, \quad (3.1)$$

the exit Mach number is found as

$$\frac{780}{300} = 1 + 0.2 M_e^2 \quad (3.2)$$

$$M_e \approx 2.83 \quad (3.3)$$

Step 2: Exit Velocity and Density

$$a_e = \sqrt{\gamma R T_e} = 347.2 \text{ m/s} \quad (3.4)$$

$$U_e = M_e a_e = 2.83 \times 347 = 982.5 \text{ m/s} \quad (3.5)$$

From the isentropic pressure ratio,

$$\frac{p^0}{p_e} = \left(1 + \frac{\gamma - 1}{2} M_e^2\right)^{\frac{\gamma}{\gamma - 1}} = 28.41 \quad (3.6)$$

$$p_e = \frac{p^0}{28.41} = 35200 \text{ Pa} \quad (3.7)$$

and thus

$$\rho_e = \frac{p_e}{R T_e} = \frac{56881}{287 \times 300} = 0.41 \text{ kg/m}^3. \quad (3.8)$$

Step 3: Wall Shear Stress and Friction Velocity

(Skin friction coefficient C_f from Prandtl's one-seventh-power law for a turbulent flat plate with experimental calibration [20])

$$C_f = 0.0592 \text{Re}_x^{-1/5} \quad \text{Re}_x = \frac{\rho_e U_e L}{\mu} \quad (3.9)$$

For $L = 0.02527 \text{ m}$

$$\text{Re}_x = 5.79 \times 10^5 \quad (3.10)$$

$$C_f = 0.004167 \quad (3.11)$$

Calculate the wall shear stress:

$$\tau_w = \frac{1}{2} C_f \rho U^2 = 824.6 \quad (3.12)$$

Calculate friction velocity:

$$u_\tau = \sqrt{\frac{\tau_w}{\rho}} = 44.85 \text{ m/s} \quad (3.13)$$

Step 4: First Cell Height

For a target $y^+ = 1$,

$$y_1 = \frac{y^+ \mu}{\rho_e u_\tau} = \frac{1.758 \times 10^{-5}}{0.41 \times 44.85} = 9.56 \times 10^{-7} \text{ m}. \quad (3.14)$$

Hence, the approximately calculated first-layer height corresponding to $y^+ = 1$ is $y_1 \approx 0.956 \mu\text{m}$. This value was used as an initial estimate for the mesh construction, i.e., y_1 was set to be at least this value. After several iterations and preliminary simulations, the optimal first-layer heights were determined to be approximately $5.9 \times 10^{-8} \text{ m}$ at the throat and $2.5 \times 10^{-7} \text{ m}$ at the inlet. Using the OpenFOAM post-processing utility `rhoCentralFoam -postProcess -func yPlus (-latestTime)`, the actual y^+ values were computed and visualised in ParaView. Figures 3.5 and 3.6 show the resulting y^+ distributions along the nozzle wall and the ground surface, respectively.





Figure 3.5: y^+ at Wall

Estimation of the Boundary-Layer Thickness at Throat

Because the most important part of the mesh is at the throat, that is the place taken into consideration. Using values of temperature ($T \approx 800K$) and Mach ($M \approx 1$) near the throat section following physical quantities can be calculated:

$$a = \sqrt{\gamma RT} = 567 \text{ m/s}, U = a \times M = 567 \text{ m/s}, p_t = \frac{p^0}{(1 + \frac{\gamma-1}{2} M^2)^{\frac{\gamma}{\gamma-1}}} = 528281 \text{ Pa},$$

$$\rho_t = 2.3 \text{ kg/m}^3 \rightarrow Re_x = 9.37 \times 10^5$$

Figure 3.6: $y+$ at Ground

An approximate estimation of the turbulent boundary-layer thickness at the nozzle throat can be obtained using this formula from empirical research into turbulent boundary layers done by Schlichting in the 1970s [22]:

$$\delta(x)_{\text{turbulent}} = 0.37 \frac{x}{\text{Re}_x^{0.2}} \approx 0.3 \text{ mm} \quad (3.15)$$

This value provides an order-of-magnitude estimate of the boundary-layer thickness expected at the nozzle throat. During the mesh generation process, several factors were taken into account: achieving a value of y_1 lower than the target, ensuring that the boundary layer thickness δ at the throat matched the calculated value, maintaining a growth ratio close to 1.2 and ensuring that the outermost boundary layer cell size smoothly transitioned to the dimensions of the external cells. The final mesh represents a balance between all of these considerations.

3.5 Thermophysical Properties of Nitrogen Gas

Nitrogen (N_2) was selected as the working fluid for all simulations, in compliance with the experiments conducted, representing an inert diatomic gas expanding through a converging-diverging nozzle into a low-pressure environment. Nitrogen was modelled as a *perfect gas* with constant thermophysical properties.

3.5.1 Equation of State and Thermodynamic Model

The perfect gas equation of state was adopted,

$$p = \rho RT, \quad (3.16)$$

where $R = R_u/M$ is the specific gas constant, with the universal gas constant $R_u = 8.3145 \text{ J/(mol K)}$ and the molar mass of nitrogen $M = 28.0134 \text{ g/mol}$. This yields a specific gas constant $R = 296.8 \text{ J/kgK}$.

A calorically perfect gas model (**hConst**) was employed, assuming a constant specific heat capacity at constant pressure c_p . It was estimated using the NASA Glenn

thermochemical data set [18]. The corresponding seven-coefficient polynomial for the temperature-dependent specific heat (in J/(mol · K)) is given by

$$\frac{c_p(T)}{R_u} = a_1 T^{-2} + a_2 T^{-1} + a_3 + a_4 T + a_5 T^2 + a_6 T^3 + a_7 T^4,$$

where R_u is the universal gas constant and the coefficients a_i for nitrogen in the temperature range $200 \text{ K} \leq T \leq 1000 \text{ K}$ are:

$$\begin{aligned} a_1 &= 2.210371497 \times 10^4, \\ a_2 &= -3.818461820 \times 10^2, \\ a_3 &= 6.082738360, \\ a_4 &= -8.530914410 \times 10^{-3}, \\ a_5 &= 1.384646189 \times 10^{-5}, \\ a_6 &= -9.625793620 \times 10^{-9}, \\ a_7 &= 2.519705809 \times 10^{-12}. \end{aligned}$$

Using this polynomial at $T = 300 \text{ K}$ yields an approximate molar specific heat of $c_p \approx 29.1 \text{ J/(mol · K)}$, or $c_p \approx 1040 \text{ J/(kg · K)}$ when normalized by the molar mass of nitrogen ($M = 28.0134 \text{ g/mol}$). As for the ratio of specific heats

$$\gamma = 1.4.$$

3.5.2 Transport Properties: Sutherland Law Coefficients for Nitrogen

The dynamic viscosity of a gas varies with temperature, and for high-fidelity simulations of compressible flows, it is important to account for this variation accurately. For nitrogen (N_2), *Sutherland law* [21] was employed, which provides a reliable model for the temperature dependence of viscosity:

$$\mu(T) = \mu_0 \frac{T_0 + T_s}{T + T_s} \left(\frac{T}{T_0} \right)^{3/2}, \quad (3.17)$$

where $\mu(T)$ is the dynamic viscosity at temperature T , μ_0 is the reference viscosity at reference temperature T_0 , and T_s is the Sutherland temperature.

In the present study, the reference temperature is chosen as $T_0 = 300 \text{ K}$, with a corresponding reference dynamic viscosity of $\mu_0 = 1.79 \times 10^{-5} \text{ Pa · s}$. The Sutherland temperature is taken as $T_s = 110.33 \text{ K}$, which is a standard value for nitrogen [30].

The Sutherland model coefficient A_s required by OpenFOAM is calculated from the reference values using:

$$A_s = \frac{\mu_0 (T_0 + T_s)}{T_0^{3/2}}. \quad (3.18)$$

Substituting the chosen reference parameters yields:

$$\begin{aligned} T_0^{3/2} &= 300^{3/2} \approx 5196.15, \\ \mu_0(T_0 + T_s) &= 1.79 \times 10^{-5} \cdot (300 + 110.33) \approx 0.007345, \\ A_s &= \frac{0.007345}{5196.15} \approx 1.48 \times 10^{-6}. \end{aligned}$$

This approach ensures that the viscosity variation with temperature is accurately captured throughout the nozzle flow.

3.5.3 Summary of Properties

Table 3.2 summarises gas propriety values used throughout the simulations and subsequent calculations.

Table 3.2: Thermophysical and transport properties of nitrogen used in the simulations.

Property	Symbol	Value
Molar weight	M	28.0134 g/mol
Specific heat capacity	c_p	1040 J/kgK
Specific heat ratio	γ	1.4
Dynamic viscosity (300 K)	μ_0	1.79×10^{-5} Pa s
Sutherland constant	A_s	1.48×10^{-6}
Sutherland temperature	T_s	110.33 K
Prandtl number	Pr	0.718
Heat of formation	h_f	0 J/kg
Boltzmann constant	k_B	1.38×10^{-23} J/K
Molecular mass	m	4.65×10^{-26} kg
Specific gas constant	R	296.8 J/(kgK)
VHS viscosity exponent	ω	0.74

3.5.4 Model Simplification Reasons

An alternative approach considered was the use of the **janaf** thermodynamic model, which provides temperature-dependent polynomial correlations for specific heat, enthalpy, and entropy. Although the **janaf** model offers greater physical accuracy across wide temperature ranges, its use in this case led to severe numerical instabilities during the initialization and early transient stages of the simulation. In particular, convergence issues were observed in the iterative evaluation of temperature from enthalpy. These difficulties typically arise when very low static pressures and temperatures are encountered (e.g., in supersonic or underexpanded jets), causing the solver to attempt extrapolation of the polynomial coefficients beyond their valid range.

Given that the flow field in the present configuration involves extreme expansions, with local static temperatures predicted between approximately 15 K and 780 K, the constant-property assumption was adopted as a more stable and computationally robust alternative.

3.6 DSMC Simulation Setup and Coupling with CFD Nozzle Flow

The external flow field downstream of the nozzle is simulated using the Direct Simulation Monte Carlo (DSMC) method, implemented via **dsmcFoam+**. A hybrid approach is adopted in which the internal nozzle flow is computed using continuum CFD, while the external plume expansion is resolved using DSMC in order to capture rarefaction and non-equilibrium effects.

The DSMC computational domain does not include the nozzle geometry, as shown in Fig. 3.7. Instead, the domain begins at the nozzle exit plane and the flow conditions at this location are obtained from the CFD solution and imposed as inlet boundary conditions for the DSMC simulation.




Figure 3.7: DSMC simulation mesh (SoH4)

The DSMC simulation is performed under the assumption of axisymmetry using the `dsmcAxisymmetric` coordinate system. This allows for a reduction in computational cost while preserving the relevant flow physics.

Radial particle weighting is employed, with a maximum weighting factor of 1000, in order to maintain a sufficient number of computational particles across the domain.

3.6.1 Variation of DSMC Parameters with Stand-Off Height

Three DSMC simulations were performed corresponding to different stand-off heights. While the physical models and boundary conditions remained unchanged, the numerical parameters were adapted in each case to balance computational cost and accuracy.

In DSMC, the spatial and temporal resolutions are ideally chosen such that the cell size is on the order of, or smaller than, the local mean free path λ , and the time step is smaller than the mean collision time τ . However, due to computational constraints, these criteria were not strictly satisfied throughout the domain.

In particular, the computational meshes employed are relatively coarse by DSMC standards and in several regions the cell size exceeds the local mean free path by a factor greater than unity. This introduces a potential source of error in the collision modelling. Nevertheless, comparison with reference DSMC simulations for the same configuration indicates that the resulting discrepancies remain within acceptable limits for the present study.

The numerical parameters were varied with stand-off height to reflect the changing flow conditions. At higher SoH, the flow becomes more rarefied, resulting in larger mean free paths and collision times. This allows for the use of larger cell sizes, larger time steps and fewer equivalent particles, thereby improving computational efficiency. Conversely, lower SoH cases require finer resolution and are therefore more computationally demanding.

An additional limitation arises from the number of particles per cell. In the higher SoH cases (8 and 16), some regions exhibit a low particle population, leading to increased

statistical noise in derived quantities such as the Mach number. In contrast, the lowest SoH case (SoH4) maintains more than five particles per cell in most regions (as indicated by the field `dsmcNMean_N2` Fig. 3.8), resulting in improved statistical convergence and smoother flow fields.




Figure 3.8: Colour map of particles per cell of SoH4

Despite these limitations, the DSMC results show good agreement with independent simulations of the same configuration. Therefore, the chosen numerical parameters are considered adequate for capturing the main physical features of the flow.

3.6.2 Gas Model and Molecular Properties

The working fluid is nitrogen (N_2), modeled as a diatomic gas. The molecular properties used in the DSMC simulation are:

- Molecular mass: $m = 4.65 \times 10^{-26}$ kg
- Molecular diameter: $d = 4.17 \times 10^{-10}$ m
- Rotational degrees of freedom: 2
- Vibrational modes: neglected

A non-reacting gas model is assumed, with no ionization or chemical reactions.

3.6.3 Collision Model

Intermolecular collisions are modeled using the Larsen–Borgnakke Variable Hard Sphere (VHS) model. This allows for energy exchange between translational and rotational modes.

The reference temperature is set to 273 K and the rotational relaxation collision number is fixed at 5. Vibrational relaxation is disabled, consistent with the assumption of negligible vibrational excitation under the simulated conditions.

Collision partner selection is performed using the no-time-counter (NTC) method.

3.6.4 Boundary Conditions

Inlet Boundary

The inlet boundary corresponds to the nozzle exit plane obtained from the CFD simulation. A field-based inflow condition (`dsmcFreeStreamInflowFieldPatch`) is used, allowing spatially varying macroscopic quantities (velocity, temperature and number density) to be imposed directly.

Particles are injected assuming a Maxwellian velocity distribution consistent with the local macroscopic properties. This implies that the flow at the inlet is assumed to be in thermodynamic equilibrium.

Far-Field Boundaries

The lateral and upper boundaries are treated as open boundaries using a combination of particle deletion and free-stream inflow conditions. The ambient conditions imposed at these boundaries are:

$$T_{\text{trans}} = 300 \text{ K} \quad (3.19)$$

$$T_{\text{rot}} = 300 \text{ K} \quad (3.20)$$

$$n_{\text{N}_2} = 4.34578 \times 10^{20} \text{ m}^{-3} \quad (3.21)$$

The free-stream velocity is set to zero.

Wall Boundaries

Solid surfaces are modeled using diffuse reflection with complete thermal accommodation. The wall temperature is prescribed as:

$$T_{\text{wall}} = 300 \text{ K} \quad (3.22)$$

Particles re-emitted from the surface are assumed to follow a Maxwellian distribution corresponding to the wall temperature.

3.6.5 Particle Representation and Sampling

The SoH4 simulation employs 3×10^8 , SoH8 employs 9×10^8 and SoH16 employs 1.3×10^8 equivalent particles ensuring adequate statistical resolution of the flow field.

Macroscopic quantities are computed using cell-based averaging via the `dsmcVolFields` model. The mean free path is also evaluated locally using a reference temperature of 273 K.

3.6.6 Coupling Strategy and Limitations

The coupling between CFD and DSMC is one-way. The CFD solution provides the inlet boundary conditions for the DSMC simulation, but no feedback from DSMC to CFD is considered.

It is assumed that the flow at the nozzle exit can be represented by a Maxwellian distribution based on the CFD-derived macroscopic properties. This assumption neglects possible non-equilibrium effects at the interface between the continuum and rarefied regions.

Furthermore, any numerical inaccuracies present in the CFD solution, such as deviations in total energy, are directly propagated into the DSMC simulation through the inlet boundary condition.

These limitations were taken into account when interpreting the DSMC results.

Chapter 4

Numerical evaluation

4.1 Grid study

With the goal of examining the mesh's independence, this simulation was conducted using varying numbers of the cells. An initial coarse mesh had been developed for interpolation purposes (for further details consult the appendix), however, its accuracy proved insufficient and it is no longer considered in the analysis. Consequently, it was necessary to verify that the new mesh provided results of adequate resolution.

The baseline solution obtained with the initial mesh was first analysed once the simulation reached a satisfactory level of convergence, with all residual imbalances falling below 1%. A second mesh was then generated with a finer discretisation throughout the domain. The refined mesh consisted of 61 647 cells in total (61 108 hexahedra and 539 prisms). Figure 4.1 shows this mesh.




Figure 4.1: Mesh used in the mesh independence Study

A second simulation was performed on this refined mesh and run to convergence. The monitored point values obtained from both meshes were then compared. If the results agreed within acceptable tolerances, the initial mesh could be considered sufficiently resolved. Otherwise, any discrepancies would indicate mesh-dependent behaviour, implying

that additional refinement would be required.

To assess the sensitivity of the solution to the spatial resolution, axial profiles of the main flow quantities were extracted along the jet centreline. Figure 4.2 to 4.5 compares the distributions of Mach number, pressure, temperature and density obtained using the two meshes considered.




Figure 4.2: Centerline distributions of Mach number




Figure 4.3: Centerline distributions of Pressure

The close agreement between the profiles indicates that the solution is weakly dependent on the grid resolution and that the main features of the under-expanded jet, including




Figure 4.4: Centerline distributions of Temperature





Figure 4.5: Centerline distributions of Density


the shock-cell structure, are captured consistently on both meshes.

4.2 CFD results

To analyse the flow structure generated by the nozzle expansion and its interaction with the impingement surface, colour maps of the main flow variables were extracted from the numerical solution. In particular, Mach number, pressure, temperature and density distributions are presented for the different stand-off heights considered. These variables provide complementary information regarding the kinematic and thermodynamic behaviour of the expanding jet.




(a) SoH = 4 cm



(b) SoH = 8 cm


Figure 4.6 shows the Mach number colour maps for the three stand-off heights SoH considered. The Mach number colour maps provide a clear visualization of the acceleration and overall structure of the expanding jet. Immediately downstream of the nozzle exit, the flow undergoes a strong Prandtl-Meyer expansion due to the large pressure difference between the nozzle exit and the surrounding near-vacuum environment. This expansion



(c) SoH = 16 cm

Figure 4.6: Colour maps of Mach number for the three stand-off heights considered.

results in a rapid increase in Mach number and the formation of a curved expansion fan originating at the nozzle lip.




(a) SoH = 4 cm


The jet subsequently continues to expand as it propagates downstream, producing a progressively larger region of highly supersonic flow.

The outer boundary of the high-Mach region corresponds to the edge of the expanding plume, where the supersonic jet interacts with the surrounding low-density environment. This boundary highlights the radial spreading of the jet as it accelerates away from the nozzle.

Near the ground surface, a shock wave is produced. In this area the jet momentum is redirected along the wall with a reduction in Mach number and the formation of a radial wall jet along the surface.



(b) SoH = 8 cm



(c) SoH = 16 cm


Figure 4.7: colour maps of pressure for the three stand-off heights considered.

For smaller stand-off heights the jet interacts with the impingement surface shortly after exiting the nozzle, while larger distances allow the jet to expand further before encountering the wall. This leads to higher local Mach numbers and a wider expansion region before the jet impinges on the surface.


The pressure colour maps corresponding to the three stand-off heights are shown in Figure 4.7. The logarithmic scale is used when mapping data to colours in order to give a clearer visual of the changes in pressure, in fact the same strategy is used for the density colour maps. Pressure colour maps are particularly useful for identifying expansion and compression regions within compressible flows. In the present case, they highlight the strong pressure drop associated with the expansion of the jet as it exits the nozzle.

In the vicinity of the ground surface, the pressure increases locally due to the deceleration of the flow as the jet approaches the wall. Higher the SoH stronger the shock wave,

which leads to a lower pressure after the shock for greater heights.



(a) SoH = 4 cm



(b) SoH = 8 cm


The temperature colour maps are presented in Figure 4.8.

As the jet expands after leaving the nozzle, the rapid acceleration of the flow results in a significant decrease in static temperature. This cooling effect is characteristic of supersonic expansions, where internal energy is converted into kinetic energy.

The lowest temperatures are therefore observed in regions where the flow reaches its highest Mach numbers.

The stand-off height influences the spatial extent of the cooling region. Larger stand-off distances allow the jet to expand over a longer distance before interacting with the wall, resulting in a more extended low-temperature region both radially and within the jet core.

Density colour maps for the three configurations are shown in Appendix C.



(c) SoH = 16 cm

Figure 4.8: colour maps of temperature for the three stand-off heights considered.

4.2.1 Observation on Total Temperature Behaviour in CFD Simulations

During the analysis of the CFD results, a non-physical increase in total temperature was observed in the divergent section of the nozzle, Fig. 4.9 and Fig. 4.10. In the present simulations, the flow is not strictly adiabatic nor inviscid due to the presence of wall heat transfer and viscous effects. Therefore, the total enthalpy is not expected to be perfectly conserved throughout the domain.

However, in the core flow region, away from solid boundaries, heat transfer and viscous dissipation are relatively limited. Under these conditions, the total enthalpy (and, under the assumption of constant specific heat capacity, the total temperature) is expected to remain approximately constant along streamlines.

The observed increase in total temperature in the divergent section therefore exceeds what can be attributed to physical effects alone.

Similar behaviour has been observed in other simulations of comparable configurations conducted within the same research group. In particular, cases performed using the `rhoPimpleFoam` solver also showed a deviation from the expected constant total temperature, although the magnitude of the increase was smaller.

In those simulations, the thermodynamic model included a temperature-dependent specific heat capacity. The resulting variation in total temperature was therefore reduced compared to the present case, where a constant specific heat capacity was assumed. This suggests that the thermodynamic modelling plays a role in the observed discrepancy.

Nevertheless, the magnitude of the total temperature increase observed in the present work cannot be attributed solely to the assumption of constant specific heat. It is more likely that the behaviour arises from a combination of factors, including the thermodynamic model and numerical effects associated with the solver. In particular, the use of a central scheme in high-Mach-number regions may introduce numerical dissipation, which can lead to artificial increases in total energy.

It is therefore concluded that the observed rise in total temperature is primarily a numerical artifact, amplified by the high-speed expansion in the divergent section and influenced by the simplifying assumptions adopted in the thermodynamic model.




Figure 4.9: Total temperature colour map for the SoH4 CFD simulation




Figure 4.10: Total temperature plot along the central axis for the SoH4 CFD simulation

A further observation can be made regarding the behaviour of the flow downstream of the shock. In the CFD simulations, although an increase in total temperature is observed in the divergent section, this discrepancy appears to diminish after the shock. The post-shock flow is characterized by relatively low velocities and the static temperature

approaches the total temperature. Consequently, the static temperature field appears physically consistent, masking the upstream increase in total temperature.

In contrast, the DSMC simulations use the flow properties at the nozzle exit as inlet conditions, where the total temperature is already affected by the aforementioned numerical increase. Since DSMC directly evolves the molecular energy without the same numerical dissipation mechanisms, the elevated total energy is preserved. After the shock, this results in a higher static temperature compared to the CFD solution, as will be observed in the following sections.

This difference highlights the sensitivity of the coupled CFD-DSMC approach to the accuracy of the inlet conditions.

4.2.2 Effect of Boundary Conditions on Pressure and Density Fields

In the CFD simulations, a region of significantly reduced pressure and density is present near the horizontal section of the domain connected to the vertical wall. This behaviour is associated with the boundary condition applied in this region, which was treated as an inlet-outlet boundary. The high-speed jet emerging from the nozzle induces a strong entrainment effect, drawing fluid from this section of the domain. As a result, a local depletion of mass occurs, leading to lower pressure and density values compared to the surrounding flow.

This effect is clearly illustrated in Fig. 4.11, where the velocity field is represented using glyphs. The vectors indicate a flow directed from the lateral boundary, confirming the presence of entrainment driven by the main jet.




Figure 4.11: Velocity field represented using glyphs in the CFD simulation.

4.3 DSMC results

The DSMC simulations capture the expansion of the supersonic jet into the surrounding environment, forming a characteristic plume structure. The Mach number contours for the three stand-off heights are shown in Fig. 4.12 - 4.14. In each figure, the upper half corresponds to the CFD solution, while the lower half shows the DSMC results. For comparison purposes, the colour maps for both CFD and DSMC results were rescaled using a common range, defined by the global minimum and maximum values across both

datasets. This ensures that identical colours correspond to the same physical values in both simulations.




Figure 4.12: Soh4 CFD and DSMC colour map of Ma




Figure 4.13: Soh8 CFD and DSMC colour map of Ma




Figure 4.14: Soh16 CFD and DSMC colour map of Ma

In all cases, the curved shock structure is observed in both CFD and DSMC simulations.

For the lowest stand-off height ($\text{SOH} = 4$), the DSMC results show good agreement with the CFD solution and the flow field appears relatively smooth due to the higher number of particles per cell.

The comparison between CFD and DSMC results reveals several important differences in the predicted plume structure. In all cases, the CFD solution exhibits a broader plume compared to the DSMC results. This difference is attributed to the reduced numerical diffusion in DSMC and its ability to more accurately capture rarefaction effects, whereas the continuum CFD approach tends to over-predict plume spreading.

While both methods capture the overall expansion behaviour and the increase in Mach number with stand-off height, the DSMC results exhibit increased statistical noise. This is due to the reduced number of particles per cell, which affects derived quantities like Mach number.

A clear difference between CFD and DSMC results is the variation in shock location with stand-off height.




Figure 4.15: SoH16 CFD shock location on the central axis




Figure 4.16: SoH16 DSMC shock location on the central axis

For the lowest stand-off height ($\text{SOH} = 4$), both methods predict the shock location with very good agreement. This indicates that, under these conditions, the flow remains sufficiently close to the continuum regime for the CFD approach to provide an accurate description of the plume expansion.

At intermediate stand-off height ($\text{SOH} = 8$), a discrepancy in shock position begins to emerge. The DSMC results predict a shock located closer to the jet core compared to the CFD solution.

This difference becomes significantly more pronounced at the highest stand-off height ($\text{SOH} = 16$), Fig. 4.15 - 4.16. In this case, the DSMC simulation shows a clearly more compact plume, with the shock located substantially closer to the nozzle axis.

A noticeable discrepancy between CFD and DSMC results is observed in the pressure and density fields, Fig. 4.17 and Fig. 4.18, particularly in the region adjacent to the left farfield. The cause of this behaviour in the CFD simulation is due to the boundary condition applied in this region, as already explained in the section 4.2.2. In contrast, the DSMC simulations do not exhibit the same level of density and pressure reduction in this region. This difference can be attributed to both the treatment of boundary conditions and the fundamentally different nature of the DSMC method. In particular, the particle-based formulation and the imposed inflow conditions limit the extent to which large-scale entrainment effects can develop in the same manner as in the continuum CFD approach.

Additional contour plots of the flow field variables for all stand-off heights are provided in Appendix C for completeness.

4.4 Centreline Comparison Between CFD and DSMC

To provide a quantitative comparison between the CFD and DSMC solutions, profiles of pressure, temperature, density and velocity magnitude were extracted along the centreline of the flow. These profiles allow for a direct assessment of the agreement between the two methods and help identify regions where deviations begin to emerge.

It should be noted that centreline profiles represent a one-dimensional slice of the flow field. As such, they are not sufficient on their own to fully characterize the spatial extent of discrepancies between the two approaches. However, they provide a useful and clear indication of the magnitude and location of differences, complementing the contour-based analysis.

For the lowest stand-off height Fig. 4.19-4.21, the profiles of pressure, density and velocity magnitude show excellent agreement, with only minor deviations.

Temperature profiles show noticeable deviations between CFD and DSMC for all stand-off heights.




Figure 4.17: Soh4 CFD and DSMC colour map of pressure field




Figure 4.18: Soh4 CFD and DSMC colour map of density field

For SoH4 and SoH8 Fig. 4.23, temperature differences exceeding 5% are observed from approximately $x = 0.02$ onwards. Although these deviations are relatively small in absolute terms, they indicate that non-equilibrium effects begin to develop. For the




Figure 4.19: Pressure along the centreline SoH4




Figure 4.20: Density along the centreline SoH4

highest stand-off height Fig.4.26 temperature differences exceeding 5% from approximately $x = 0.05$ onwards.

In addition to temperature, differences in other variables become more evident at higher stand-off heights. For SoH16, the shock location predicted by CFD and DSMC differs significantly. Upstream of the shock, the velocity magnitude Fig.4.25 shows deviations but these are negligible (max $\approx 3\%$), while density Fig.4.24 differences exceed 5% from approximately $x = 0.0755$ onwards.

The remaining plots can be found in the Appendix C




Figure 4.21: Velocity along the centreline SoH4




Figure 4.22: Temperature along the centreline SoH4

4.5 Ground Pressure Comparison with Experimental Data

The pressure distribution from both solvers along the ground was analysed and compared with available experimental measurements. This provides an additional reference for assessing the accuracy of the numerical simulations. The coordinate $x = 0.0201m$ represents the radial direction along the ground, with $x = 0$ corresponding to the centreline of the jet impingement region.

It should be noted that the experimental measurements are affected by significant




Figure 4.23: Temperature along the centreline SoH8




Figure 4.24: Density along the centreline SoH16

uncertainty, with all reported values having an estimated uncertainty of at least ± 700 Pa. As a result, the comparison presented here should be interpreted qualitatively rather than quantitatively.

The available experimental data for the three stand-off heights are summarised in Table 4.1.

The two numerical solutions show good agreement over most of the radial extent, with significant deviations confined to a region near the centreline which increases in size with stand-off height.




Figure 4.25: Velocity along the centreline SoH16




Figure 4.26: Temperature along the centreline SoH16

4.6 Mean Free Path Estimation

The molecular mean free path represents the average distance travelled by a molecule between two successive collisions. It is a key parameter in rarefied gas dynamics since it directly appears in the definition of almost all parameters considered and therefore determines the validity of the continuum hypothesis. Depending on the available flow quantities and the modelling approach adopted, the mean free path can be evaluated using several equivalent formulations derived from kinetic theory.

Table 4.1: Experimental ground pressure measurements at different stand-off heights. All values are subject to an uncertainty of at least ± 700 Pa.

Case	x_1 (m)	P_1 (Pa)	x_2 (m)	P_2 (Pa)	Notes
SoH4	0	4900	2.01×10^{-2}	430	–
SoH8.95	0	2830	2.01×10^{-2}	391	Low reliability
SoH16	0	490	2.01×10^{-2}	324	–




Figure 4.27: Ground pressure distribution along the radial direction for the SoH4




Figure 4.28: Ground pressure distribution along the radial direction for the SoH8




Figure 4.29: Ground pressure distribution along the radial direction for the SoH16

Kinetic theory formulation

In classical kinetic theory, assuming a dilute gas composed of hard-sphere molecules, the mean free path can be expressed as

$$\lambda = \frac{1}{\sqrt{2} n \pi d^2}$$

where n is the molecular number density and d is the molecular diameter. The factor $\sqrt{2}$ accounts for the relative motion of colliding molecules. Using the ideal gas relation $p = nk_B T$, where k_B is the Boltzmann constant, the expression can be rewritten in terms of macroscopic variables as

$$\lambda = \frac{k_B T}{\sqrt{2} \pi d^2 p}$$

where p is the pressure and T is the temperature. When this formulation is used, the value of d must represent an effective hard-sphere collision diameter. For nitrogen, value is close to $d \approx 4.1 \times 10^{-10}$ m.

Viscosity-based formulation

When molecular properties such as the collision diameter are not explicitly available, the mean free path can be expressed using macroscopic transport properties. By combining the kinetic-theory expression for viscosity with the ideal gas law, the mean free path can be written as

$$\lambda = \frac{\mu}{p} \sqrt{\frac{\pi R T}{2}}$$

where μ is the dynamic viscosity, evaluated using Sutherland's law, and R is the specific gas constant.

This formulation is particularly convenient in continuum simulations since pressure, temperature and viscosity are readily available from CFD solutions.

DSMC reference diameter formulation

In Direct Simulation Monte Carlo (DSMC) methods, the collision dynamics are defined through a molecular diameter consistent with the chosen collision model. In the Variable Hard Sphere (VHS) model introduced by Bird [4], the reference diameter is determined from the viscosity at a reference temperature as

$$d_{ref} = \left(\frac{5}{16} \frac{\sqrt{\pi m k_B T_{ref}}}{\mu_{ref}} \frac{1}{\Gamma(2.5 - \omega)} \right)^{1/2}$$

where m is the molecular mass, μ_{ref} is the viscosity at the reference temperature T_{ref} , ω is the viscosity-temperature exponent, and $\Gamma(\cdot)$ denotes the Gamma function, which extends the factorial function to non-integer arguments and is defined as

$$\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt.$$

The Gamma function arises from the analytical evaluation of the collision integrals appearing in the Chapman–Enskog solution of the Boltzmann equation when deriving the viscosity of gases with velocity-dependent collision cross sections. For nitrogen, a commonly used value of the viscosity exponent is $\omega \approx 0.74$, which yields $\Gamma(2.5 - \omega) \approx 0.92$. Substituting the corresponding gas properties results in a reference molecular diameter of approximately $d_{ref} \approx 3.7 \times 10^{-10}$ m, which is widely used in DSMC simulations of nitrogen flows. Because the VHS diameter is a model parameter rather than a true hard-sphere diameter, substituting it directly into the hard-sphere mean free path expression leads to small discrepancies.

Although the different formulations arise from distinct derivations, they are physically equivalent and represent the same molecular transport scale. The kinetic-theory form is most commonly used when molecular parameters are known, while the viscosity-based expression is generally preferred in continuum simulations. In DSMC simulations, the molecular diameter is typically obtained from the VHS or VSS collision models.

In the present work the viscosity-based expression is adopted. This choice is supported by a comparison with the mean free path obtained from the DSMC simulation, which showed better agreement with the viscosity-based estimate than with the hard-sphere formulation using the VHS reference diameter.

4.7 Knudsen Number

In the present study, the characteristic length L is taken as the nozzle throat diameter, 1 mm, which provides a representative scale for the rarefaction effects in the flow.

The Knudsen number, Kn , was used to identify regions where the continuum assumption may no longer be valid. It, along with all the other parameters, was calculated in post-processing in paraview. In order to visualise where the values between the simulations begin to differ, different measures were used. These include observing plots over a line of different variables as well as observing contours of these variables on the geometry itself. In the literature, values between 0.01 and 0.05 are typically considered as the limit for Kn .




Figure 4.30: SoH4 with $Kn = 0.05$





Figure 4.31: SoH8 with $Kn = 0.05$

Figures 4.30 - 4.32 show the contour $Kn = 0.05$, with the mach field colour map underneath, which could be used as the interface for a coupled CFD-DSMC simulation. Appendix C contains figures with this contour with a grid showing the exact locations of breakdown.

The criterion is relatively conservative, which may lead to larger DSMC regions in a coupled approach, increasing computational cost while improving solution robustness.

Figure 4.32: SoH16 with $Kn = 0.05$

4.8 Bird’s Parameter

Bird’s breakdown parameter was primarily evaluated using the directional derivative along the local velocity vector, rather than the conventional gradient magnitude used in the definition of Kn_{GLL} . This formulation better represents the dominant direction of molecular transport and helps avoid overestimating continuum breakdown in strongly anisotropic flow fields, such as those found in supersonic expansions. The P_{max} parameter which was presented in sec.2.4.2 is not reproduced here because it presented no significant differences compared to the original P parameter. While P calculated using gradients of variables other than density proved to be no more accurate, so these are also not presented here.

As mentioned before, a value of $P = 0.05$ is generally adopted, Figures 4.33 - 4.35 show the contours with this value.

The results indicate that this parameter predicts the flow to remain largely within the continuum regime upstream of the shock. This prediction is reasonably accurate for the lowest stand-off height case, however, for the higher stand-off heights, the breakdown parameter fails to accurately capture the onset of non-equilibrium effects.

In particular, the shock is predicted to form by DSMC in regions where the breakdown parameter still indicates continuum conditions. This suggests that the parameter underestimates the extent of non-equilibrium effects in high-Mach-number regions, especially in the vicinity of strong gradients such as shock waves.

On the other hand, the breakdown parameter seems to provide a more reliable indication of rarefaction effects along the plume boundaries, where gradients are less abrupt.

When a lower threshold value (e.g., $P = 0.02$) is considered, localized regions begin to emerge that correlate with discrepancies between CFD and DSMC results. Figure 4.36 shows a comparison of the pressure contours on SoH16, which shows discrepancies between CFD and DSMC that extend over a finite region of the flow. In these areas, differences on the order of 10-20% are observed, indicating the onset of significant deviations between the two methods.




Figure 4.33: SoH4 with $P = 0.05$






Figure 4.34: SoH8 with $P = 0.05$

Despite this improved sensitivity, the resulting regions are highly non-uniform and do not form a smooth, continuous interface. This limits the practical applicability of the parameter for defining a coupling boundary, as it does not provide a clear separation.

Figure 4.35: SoH16 with $P = 0.05$ Figure 4.36: Contour plots of $P = 0.02$ (blue), pressure = 5 Pa calculated with CFD (red) and DCMS (yellow)


4.9 Gradient Local Length Knudsen number

The gradient-based local Knudsen number, Kn_{GLL} , was evaluated using gradients of temperature, density and velocity magnitude. The maximum value obtained from these quantities was used as the representative breakdown parameter at each location in the domain.

A threshold value of $Kn_{GLL} = 0.05$ was adopted, consistent with commonly used criteria in the literature to identify the breakdown of the continuum assumption. Fig.4.37 shows the contour for the 16 cm stand of height, rest of the figures can be viewed in the Appendix C.


Kn_{GLL} provides a good indication of rarefaction effects along the plume boundaries.

As discussed in the theoretical section, Kn_{GLL} is closely related to the parameter P_Q , and therefore it is expected that both criteria exhibit similar spatial distributions. This is


 Figure 4.37: SoH16 with $K_{GLL} = 0.05$

confirmed by the present results.

Kn_{GLL} does not accurately predict the location of the shock. The shock provided by DSMC occurs in regions where the breakdown parameter remains below the adopted critical threshold.


 Figure 4.38: Scaled colour map of Kn_{GLL} for SoH16

Furthermore, Kn_{GLL} takes on very low values within the core of the expanding jet, including in the vicinity of the shock, Fig.4.38. As a result, greatly reducing the threshold value (by more than an order of magnitude) moves the interface upstream, predicting more

accurately the rarefaction along the centre but overestimating it radially.

This indicates that the breakdown parameter, while useful as a general guideline, should be applied with caution and may require case-dependent interpretation.

4.10 Deborah number

The Deborah number compares the characteristic molecular relaxation time with the macroscopic flow timescale

$$De = \frac{\tau_r}{t_f}. \quad (4.1)$$

The relaxation time required for the evaluation of the Deborah number is approximated using the molecular collision time, i.e.

$$\tau_{relax} \approx \tau_c \approx \frac{\lambda}{\bar{c}}.$$

where λ is the mean free path and $\bar{c}$ is the mean molecular speed. This approximation is commonly adopted for simple gases and provides a reasonable estimate of the characteristic time scale for momentum and energy equilibration. However, it should be noted that relaxation processes may involve multiple collisions, particularly in strongly non-equilibrium regions such as shock waves and therefore this approximation introduces some uncertainty in the evaluation of the Deborah number.




Figure 4.39: SoH16 Contour of $De = 0.35$, with a logarithmic pressure colour map underneath

Assuming a Maxwellian velocity distribution, $\bar{c}$ can be expressed as

$$\bar{c} = \sqrt{\frac{8RT}{\pi}}. \quad (4.2)$$

The macroscopic flow timescale is estimated from the ratio between a characteristic length scale and the local velocity magnitude,

$$\tau_f \approx \frac{L}{U}. \quad (4.3)$$

the Deborah number can be estimated as

$$De \approx \frac{\lambda U}{\bar{c} L}. \quad (4.4)$$

A threshold value of $De = 0.1$ is commonly used in the literature. In the present work, a higher value of $De = 0.35$ was observed to maintain validity, Fig. 4.39.

As expected from Eq. 4.4, the Deborah number exhibits behaviour similar to a rescaled Knudsen number. While this provides a consistent indication of rarefaction, it also leads to a relatively conservative estimate of the breakdown region. Fig. C.19 shows the two parameters together.

It is important to note that the Knudsen number and the Deborah number do not fail in identifying the shock region, rather they provide conservative estimates of the breakdown interface. In both cases, the predicted breakdown region extends upstream of the shock, ensuring that the non-equilibrium region is fully encompassed. While this behaviour guarantees physical accuracy, based on the threshold chosen, it may lead to an unnecessarily large DSMC region in a coupled simulation, with increased computational cost.

Chapter 5

Conclusions

The present work focused on the assessment of continuum breakdown parameters for high-speed plume expansion in near-vacuum conditions and its interaction with the ground, with the objective of evaluating their ability to identify regions where the Navier–Stokes equations are no longer valid.

A comparative analysis of several breakdown parameters highlights significant differences in their ability to identify the onset of non-equilibrium effects. At low stand-off height, where breakdown parameters remain small throughout most of the domain, good agreement between CFD and DSMC is observed. As the stand-off height increases, regions identified as near or beyond breakdown expand and discrepancies between CFD and DSMC become more pronounced. These include differences in plume spreading and shock location, providing physical evidence that supports the trends indicated by the breakdown parameters.

The Knudsen number and the Deborah number provide conservative estimates of the breakdown region, consistently predicting the interface upstream of the shock and therefore ensuring that non-equilibrium regions are fully encompassed.

In contrast, gradient-based parameters such as Bird’s parameter P and Gradient Local Length Knudsen number Kn_{GLL} offer a more localized description of the flow and are effective in identifying rarefaction along the plume boundaries. However, when using commonly adopted threshold values, these parameters fail to capture the shock location, indicating an underestimation of non-equilibrium effects in strongly compressible regions.

A reduction of the threshold value for P improves its sensitivity, leading to localized regions that correlate with observed discrepancies between CFD and DSMC results. Nevertheless, these regions are highly non-uniform and do not form a smooth interface, limiting their practical applicability for coupling purposes. While the reduction of the threshold value for Kn_{GLL} can lead to a smoother but more conservative interface.

Additional effects, including sensitivity to boundary condition treatment and anomalies in the temperature field, were identified and shown to influence the CFD results.

Overall, the results demonstrate that no single breakdown parameter is capable of simultaneously providing an accurate and robust definition of the continuum breakdown interface. Conservative criteria ensure physical reliability but at the expense of computational efficiency, while more sensitive parameters capture local deviations but lack robustness.

While the present work employs a one-way coupling between CFD and DSMC simulations, future work may focus on the development of fully coupled hybrid approaches, where the interface is dynamically determined and two-way interaction between the solvers is

taken into account. In particular, the development of adaptive or multi-parameter approaches could provide a more reliable identification of non-equilibrium regions. From a modelling perspective, the use of more advanced thermodynamic models, such as temperature dependent specific heat capacities, may help reduce discrepancies observed in the total temperature field. Finally, the influence of boundary condition treatment and numerical setup on the prediction of rarefied flows could be investigated in more detail to further improve the reliability of continuum simulations in transitional regimes.

Bibliography

- [1] pygimli mesh interpolation. https://www.pygimli.org/_tutorials_auto/1_mesh/plot_5-mesh_interpolation.html.
- [2] Hailin Bi, Yicong Zhang, Ziyang He, Guizhong Zuo, Bin Cao, Jun Zhang, Jun Wu, Qing Cao, and Xudi Wang. A coupled ns-dsmc method applied to supersonic molecular beam and experimental validation. *Vacuum*, 214:112228, 2023.
- [3] G. A. BIRD. Breakdown of translational and rotational equilibrium in gaseous expansions. *AIAA Journal*, 8(11):1998–2003, 1970.
- [4] Graeme A Bird. *Molecular gas dynamics and the direct simulation of gas flows*. Oxford university press, 1994.
- [5] Israel Borges Sebastiao. *Ab-initio Collision Models for DSMC and Their Applications to Reacting Flows*. PhD thesis, 08 2017.
- [6] Iain D Boyd, Gang Chen, and Graham V Candler. Predicting failure of the continuum fluid equations in transitional hypersonic flows. *Physics of fluids*, 7(1):210–219, 1995.
- [7] Iain D Boyd. Analysis of rotational nonequilibrium in standing shock waves of nitrogen. *AIAA journal*, 28(11):1997–1999, 1990.
- [8] S. Chapman and T.G. Cowling. *The Mathematical Theory of Non-uniform Gases: An Account of the Kinetic Theory of Viscosity, Thermal Conduction and Diffusion in Gases*. Cambridge Mathematical Library. Cambridge University Press, 1990.
- [9] Po-Heng Chen, Iain Boyd, and José Camberos. Assessment of entropy generation rate as a predictor of continuum breakdown. In *36th AIAA thermophysics conference*, page 3783, 2003.
- [10] H. Ebrahimgol, M. Aghaie Moghanloo, and A. Zolfaghari. Coupling molecular and continuum numerical methods for evaluation of thermal and mechanical drives on high rotating isotopic gas flow. *Physics of Fluids*, 36(12):126134, 12 2024.
- [11] Mohamed Gad-el Hak. The fluid mechanics of microdevices-the freeman scholar lecture. *Journal of Fluids Engineering*, 121(1):5–33, 03 1999.
- [12] Alejandro L Garcia, John B Bell, William Y Crutchfield, and Berni J Alder. Adaptive mesh and algorithm refinement using direct simulation monte carlo. *Journal of computational Physics*, 154(1):134–155, 1999.
- [13] Valiyollah Ghazanfari, Ali Akbar Salehi, Ali Reza Keshtkar, Mohammad Mahdi Shadman, and Mohammad Hossein Askari. Investigation of the continuum-rarefied flow and isotope separation using a hybrid cfd-dsmc simulation for uf6 in a gas centrifuge. *Annals of Nuclear Energy*, 152:107985, 2021.
- [14] Christopher Greenshields and Henry Weller. *Notes on Computational Fluid Dynamics: General Principles*. CFD Direct Ltd, Reading, UK, 2022.
- [15] Alexander Kurganov and Eitan Tadmor. A new semidiscrete central-upwind scheme for hyperbolic conservation laws and hamilton–jacobi equations. *SIAM Journal on Scientific Computing*, 22:1461–1488, 2000.
- [16] C David Levermore, William J Morokoff, and Balasubramanya T Nadiga. Moment

- realizability and the validity of the navier–stokes equations for rarefied gas dynamics. *Physics of Fluids*, 10(12):3214–3226, 1998.
- [17] Duncan A Lockerby, Jason M Reese, and Henning Struchtrup. Switching criteria for hybrid rarefied gas flow.
 - [18] Bonnie J McBride. *NASA Glenn coefficients for calculating thermodynamic properties of individual species*. National Aeronautics and Space Administration, John H. Glenn Research Center, 2002.
 - [19] OpenFOAM Foundation / OpenCFD. Openfoam v2112 user guide. <https://www.openfoam.com/documentation/guides/latest/doc/guide-applications-solvers-compressible-rhoCentralFoam.html>.
 - [20] Hermann Schlichting and Klaus Gersten. *Boundary-layer theory*. springer, 2016.
 - [21] William Sutherland. Lii. the viscosity of gases and molecular force. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 36(223):507–531, 1893.
 - [22] Tom-Robin Teschner. The y-plus (y+) value. <https://cfd.university/ideas/y-plus-y-values-and-its-importance-for-mesh-generation/>.
 - [23] S Tiwari. Coupling of the boltzmann and euler equations with automatic domain decomposition. *Journal of Computational Physics*, 144(2):710–726, 1998.
 - [24] Hsue-Shen Tsien. Superaerodynamics, mechanics of rarefied gases. *Journal of the Aeronautical Sciences*, 13(12):653–664, 1946.
 - [25] H. K. Versteeg and W. Malalasekera. *An Introduction to Computational Fluid Dynamics: The Finite Volume Method*. Pearson Education Limited, 2007.
 - [26] Pradeep Vijayakumar, Quanhua Sun, and Iain D Boyd. Vibrational–translational energy exchange models for the direct simulation monte carlo method. *Physics of Fluids*, 11(8):2117–2126, 1999.
 - [27] Wolfgang Wagner. A convergence proof for bird’s direct simulation monte carlo method for the boltzmann equation. *Journal of Statistical Physics*, 66(3-4):1011 – 1044, 1992.
 - [28] Wen-Lan Wang and Iain D. Boyd. Predicting continuum breakdown in hypersonic viscous flows. *Physics of Fluids*, 15(1):91–100, 01 2003.
 - [29] C. White, M.K. Borg, T.J. Scanlon, S.M. Longshaw, B. John, D.R. Emerson, and J.M. Reese. dsmcfoam+: An openfoam based direct simulation monte carlo solver. *Computer Physics Communications*, 224:22–43, 2018.
 - [30] Frank M White and Joseph Majdalani. *Viscous fluid flow*, volume 3. McGraw-Hill New York, 2006.
 - [31] 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.
 - [32] Jun Zhang, Benzi John, Marcel Pfeiffer, Fei Fei, and Dongsheng Wen. Particle-based hybrid and multiscale methods for nonequilibrium gas flows. *Advances in Aerodynamics*, 1(1):12, 2019.

Appendix A

Interpolating fields using MapFields

To reduce the computational cost of running simulations on a high-resolution mesh, a useful approach is to use the `mapFields` utility in OpenFOAM. This tool maps one or more fields relating to a given geometry onto the corresponding fields for another geometry, Fig. A.1. It is completely generalised in so much as there does not need to be any similarity between the geometries to which the fields relate. However, for cases where the geometries are consistent, which is the case in this study, `mapFields` can be executed with a special option that simplifies the mapping process [14].




Figure A.1: Interpolate data to different meshes. Reproduced from [1]

The data is mapped from the *source* to the *target*. The fields are deemed *consistent* if the geometry and boundary types, or conditions, of both source and target fields are identical. The field data that mapFields maps are those fields within the time directory specified by startFrom/startTime in the controlDict of the target case. The data is read from the equivalent time directory of the source case, unless otherwise specified and mapped onto the equivalent time directory of the target case [14].

In this work, the simulation was first carried out on a coarse mesh to obtain an initial solution. The results from the final time step of this simulation were then interpolated onto a finer mesh using mapFields. This interpolation served as the initial condition for the finer mesh simulation (0/ of the new case), helping to accelerate convergence and reduce total computation time by starting from a physically realistic state rather than default initial values.

In order to simplify the mapping process **-consistent** command line option was used when possible. Considering that it can stop working for many "simple" reasons in some simulations another method was used. In such cases mapFieldDict directory was created in the system directory of the target cases, as shown in the listing presented below. It should be noted that patchMap is not necessary if all the mappings are with identical names, as is the case for the example below, since the patches with identical names are auto-mapped in v2412. The file presented was adapted from public codes present at github.com .

Listing A.1: OpenFOAM mapFieldsDict for patch mapping

```

/*-----*/
|=====|
|  \ \   /F ield      | OpenFOAM: The Open Source CFD Toolbox |
|  \ \   /O peration  | Version:   2412                      |
|  \ \   /A nd         | Web:      www.openfoam.com           |
|  \ \   /M anipulation|                                     |
/*-----*/

FoamFile
{
    version      2.0;
    format       ascii;
    class        dictionary;
    object       mapFieldsDict;
}

// This dictionary specifies how fields are mapped from a
//source case to a target case using the mapFields utility.

// There are three different options:
// - patch exists in the source case: specify mapping (patchMap)
// - patch should be interpolated from internal values in source
//   case (cuttingPatches)
// - patch should not be mapped. Default if not in patchMap
//   or cuttingPatches

// List of pairs of target/source patches for mapping
patchMap
(

```

```
(INLET INLET)
(SYM_FRONT SYM_FRONT)
(SYM_BACK SYM_BACK)
(FARFIELD_RIGHT FARFIELD_RIGHT)
(FARFIELD_TOP FARFIELD_TOP)
(FARFIELD_LEFT FARFIELD_LEFT)
(WALL WALL)
);

cuttingPatches
(
);
```

To map these cases, the command executed was: `mapFields /pathToSourceCase -sourceTime latestTime -mapMethod cellPointInterpolate`

In this command:

- `/pathToSourceCase` specifies the directory of the source (coarse) simulation from which fields are to be mapped.
- `-sourceTime latestTime` instructs OpenFOAM to take the most recent time-step available in the source case for mapping.
- `-mapMethod cellPointInterpolate` selects the interpolation method used to transfer field values from the source mesh to the target mesh. The `cellPointInterpolate` method interpolates field values based on cell-centre points, which is suitable for mapping from a coarse mesh to a finer, possibly non-uniform mesh while preserving gradients and boundary consistency.

Although not strictly required, it is good practice to copy the `0` directory from the source case to the target case. In some instances, `mapFields` may fail to create these files automatically, which can result in the utility running without actually mapping any field values.

Considerations for the time-varying b.c.:

Special attention must be paid to boundary conditions in cases with time-varying conditions like the present study case setup. In this case, the inlet was defined with a linearly varying boundary condition for pressure. When `mapFields` is executed, the contents of the `latestTime` directory (carried over from the original `0` directory), including boundary condition files, are copied to the new case directory. As a result, the inlet boundary condition file in the finer mesh case retains the original definition, which is no longer appropriate when restarting from a later physical time.

To avoid resetting the inlet condition to its initial value from the coarse mesh simulation, this boundary condition must be manually updated, which in this case study means both the previously time-varying p^0 and static p value for the first time step indicated in `uniformTotalPressure` type inlet. Otherwise, the new simulation may incorrectly reapply the original ramped condition, leading to nonphysical behaviour or inconsistencies during the fine mesh simulation.

Wall Functions on coarser meshes:

The $k-\omega$ SST model was used with the standard low-Re wall treatment (`nutLowReWallFunction`, `kLowReWallFunction`, and `omegaWallFunction`), which allows a smooth transition between wall-resolved behavior at low y^+ and wall-function behaviour on coarser meshes, ensuring numerical robustness during mesh refinement.

Appendix B

Turbulence Boundary Conditions

B.1 Full 0/k files

```
/*-----*-- C++ --*-----*\
=====
\\      /  F i e l d      |  OpenFOAM: The Open Source CFD Toolbox
\\      /  O p e r a t i o n |  Website:  https://openfoam.org
\\      /  A n d             |  Version:   6
\\      /  M a n i p u l a t i o n |
\*-----*-----*/
FoamFile
{
    version      2.0;
    format       ascii;
    class        volScalarField;
    object       k;
}
// * * * * *

dimensions      [0 2 -2 0 0 0 0];

internalField    uniform 100;

boundaryField
{
    FARFIELD_LEFT
    {
        type      inletOutlet;
        inletValue $internalField;
        value      $internalField;
    }
    FARFIELD_TOP
    {
        type      inletOutlet;
        inletValue $internalField;
        value      $internalField;
    }
}
```

```

FARFIELD_RIGHT
{
    type            kLowReWallFunction;
    value            uniform 1e-10;
}
INLET
{
    type            turbulentIntensityKineticEnergyInlet;
    intensity        0.01;                // 1% turbulence
    value            uniform 1;
}
WALL
{
    type            kLowReWallFunction;
    value            uniform 1e-10;
}
SYM_FRONT
{
    type            wedge;
}
SYM_BACK
{
    type            wedge;
}
}
// * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * //

```

B.2 Full $0/\omega$ file

```

// * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * //
dimensions        [0 0 -1 0 0 0 0];

internalField      uniform 300000;

boundaryField
{
    FARFIELD_LEFT
    {
        type            inletOutlet;
        inletValue        $internalField;
        value            $internalField;
    }
    FARFIELD_TOP
    {
        type            inletOutlet;
        inletValue        $internalField;
        value            $internalField;
    }
}

```

```
FARFIELD_RIGHT
{
    type          omegaWallFunction;
    value         uniform 2e+10;
}
INLET
{
    type          turbulentMixingLengthFrequencyInlet;
    mixingLength  0.000567; //0.07*d_h=0.000567 [m]
    phi          phi;
    k            k;
    value        uniform 1200;
}
WALL
{
    type          omegaWallFunction;
    value         uniform 2e+10;
}
SYM_FRONT
{
    type          wedge;
}
SYM_BACK
{
    type          wedge;
}
}
// ***** //
```


Appendix C

CFD and DSMC results and Breakdown contours




Figure C.1: colour maps of density for $\text{SoH} = 4 \text{ cm}$




Figure C.2: colour maps of density for SoH = 8 cm




Figure C.3: colour maps of density for SoH = 16 cm




Figure C.4: Soh8 CFD and DSMC colour map of pressure field




Figure C.5: Soh16 CFD and DSMC colour map of pressure field




Figure C.6: Soh8 CFD and DSMC colour map of density field




Figure C.7: Soh16 CFD and DSMC colour map of density field




Figure C.8: Soh4 CFD and DSMC colour map of temperature field




Figure C.9: Soh8 CFD and DSMC colour map of temperature field




Figure C.10: Soh16 CFD and DSMC colour map of temperature field




Figure C.11: Pressure along the centreline SoH8




Figure C.12: Density along the centreline SoH8




Figure C.13: Velocity along the centreline SoH8




Figure C.14: Pressure along the centreline SoH16




Figure C.15: SoH8 with $Kn = 0.05$




Figure C.16: SoH16 with $Kn = 0.05$




Figure C.17: SoH4 with $K_{GLL} = 0.05$




Figure C.18: SoH8 with $K_{GLL} = 0.05$




Figure C.19: SoH16 with $De = 0.3$ (blue) and $Kn = 0.05$ (red)