



Vaasan yliopisto
UNIVERSITY OF VAASA

Ali Daryab Qasemi

**Development and Validation of an Automated
Numerical Flow Bench Workflow for Rapid
Prototyping of Next Generation Combustion
Systems**

Wärtsilä Finland Oy

School of Technology and Innovations
Master's Thesis in Smart Energy
Master of Science in Technology

Vaasa 2025

UNIVERSITY OF VAASA**School of School of Technology and Innovations**

Author:	Ali Daryab Qasemi		
Title of the thesis:	Development and Validation of an Automated Numerical Flow Bench Workflow for Rapid Prototyping of Next Generation Combustion Systems		
Degree:	Master of Science in Technology		
Degree Programme:	Smart Energy		
Supervisor:	Prof. Maciej Mikulski		
Instructor:	Dr. Mahmoud Gadalla		
Year:	2025	Pages:	160

ABSTRACT:

The development of next-generation combustion systems for marine applications, driven by the urgent need for decarbonization and enhanced efficiency, is often constrained by time-intensive design-validation cycles. While Computational Fluid Dynamics (CFD) is a powerful tool for analysing in-cylinder flow, its manual setup and long runtimes create a significant bottleneck for rapid prototyping.

This thesis addresses this challenge by developing and validating a robust, automated workflow for numerical flow bench simulations, specifically applied to the gas exchange system of a Wärtsilä 31 engine conducted through a flow bench test rig. The core of this work is the PyFlowBench package, a Python toolkit designed to automate the entire CFD pipeline, from parametric mesh generation and case setup using experimental data to simulation execution and post-processing. An extensive series of Reynolds-Averaged Navier-Stokes (RANS) simulations were performed using this workflow and rigorously benchmarked against a comprehensive experimental dataset.

The results demonstrate two primary achievements. First, the automated workflow successfully reduced simulation turnaround time, enabling consistent and scalable analysis. Second, the validated CFD model showed strong agreement with experimental data, particularly in high-flow regimes where measurement uncertainty was minimal, with mass flow rate predictions generally falling under a $\pm 5\%$ deviation. Crucially, the simulations also served a powerful diagnostic role, identifying and quantifying significant experimental uncertainties stemming from sensor placement and instrumentation limitations. The study also identified the need for further uncertainty quantification of experimental measurements and numerical modelling. Further, it was pinpointed that the adiabatic wall assumption was a key source of temperature prediction errors.

Ultimately, this work delivers a validated and automated platform that significantly accelerates the design-analysis cycle. It lays the essential groundwork for future advanced studies, including Large Eddy Simulations (LES), Conjugate Heat Transfer (CHT), and the development of Machine Learning surrogate models, paving the way for the rapid optimisation of more efficient and cleaner combustion systems.

KEYWORDS: Computational Fluid Dynamics (CFD), Numerical Flow Bench, Automated Workflow, Internal Combustion Engine (ICE), Gas Exchange, Validation and Verification (V&V), Rapid Prototyping, Reynolds-Averaged Navier-Stokes (RANS), Mass flow measurement, Turbulence modelling, Uncertainty analysis.

Acknowledgements

The completion of this master's thesis represents the culmination of a challenging and rewarding journey, one that would not have been possible without the support of many individuals. I would like to extend my gratitude to all of them.

I wish to express my profound gratitude to my supervisor at the University of Vaasa, Prof. Dr. Maciej Mikulski. His unwavering support and meticulous, line-by-line feedback were instrumental in shaping this thesis into its final form. His consistent input speaks volumes of his commitment to my academic development. I am also deeply indebted to my instructor at Wärtsilä, Dr. Mahmoud Gadalla. His immense patience and step-by-step guidance were the bedrock of my technical work. This project would have been an insurmountable task without his expert mentorship. Sincere thanks to Mr. Ben Smulter, whose consistent feedback provided a critical link between the experimental data and numerical results that helped in identifying key inaccuracies.

I am particularly grateful to Mr. Magu Raam Prasaad for his foundational contribution in developing the essential scripts for the automated SnappyHexMesh generation, a cornerstone of the entire workflow. I also thank the wider Wärtsilä CFD team: Mr. Eric Lendormy for entrusting me with this responsibility and for his insightful feedback; Dr. Karri Keskinen for his expert advice; and Mr. Sakari Leponiemi for his contributions to the meshing workflow and for refining the numerical case setup.

Finally, I would like to express my most heartfelt gratitude to my family, especially my parents, whose unwavering love, belief, and sacrifices have made it possible to get an exceptional education and opportunities I would not have had otherwise. Thanks to them, I have been able to pursue my goals without a worry, and this achievement is as much theirs as it is mine.

Contents

1	Introduction	16
1.1	Thesis Goals and Research Questions	21
1.2	Thesis Structure	22
2	Background and Literature Review	25
2.1	Principles of Flow Bench Measurements	25
2.1.1	Typical components of a Flow Bench System	26
2.2	Literature Survey: Valve Flow Studies and Flow Bench Applications	29
2.2.1	Fundamental Studies on Valve Flow Research	29
2.2.2	Flow Bench in Experimental Studies	43
2.2.3	Flow Bench in Numerical Studies	48
2.2.4	Flow Bench in Hybrid Experimental-Numerical Studies	53
2.3	Governing Conservation Equations	57
2.3.1	Mass	57
2.3.2	Momentum	58
2.3.3	Energy	58
2.3.4	Scalar Quantities	59
2.3.5	Equation of State	60
2.4	Overview of Turbulence Modelling	60
2.4.1	RANS Models and Equations	63
2.5	Numerical Discretisation Approaches	65
2.5.1	Finite Difference Method (FDM)	66
2.5.2	Finite Volume Method (FVM)	67
2.5.3	Finite Element Method (FEM)	69
3	Methodology	71
3.1	Experimental Setup and the Source of Data	71
3.1.1	Physical Setup	72
3.1.2	Design of Experiments	75
3.2	Numerical Setup	81
3.2.1	Computational Mesh Generation	81

3.2.2	Case Setup and Input/Output Structure	85
3.2.3	Numerical Uncertainty Analysis	94
3.3	Development of PyFlowBench Package	100
4	Results and Discussion	104
4.1	Initial Validation and Emerging Experimental Uncertainties	104
4.2	Intake Flow Configuration	107
4.2.1	RANS Turbulence Models Sensitivity Analysis on Mass Flow Rate	108
4.2.2	Effects of Pressure Ratio on Mass Flow Rate	112
4.2.3	Effects of Valve Lifts on Mass Flow Rate	114
4.3	Exhaust Flow Configuration	121
4.4	Comparison of Different Cylinder Head Designs	123
4.4.1	Cylinder Head 3 vs. Cylinder Head 1	124
4.4.2	Cylinder Head 1 Individual Intake Valves and Exhaust Valves	126
4.5	Key Findings for Experimental and Numerical Discrepancies	129
4.5.1	Influence of Absolute Pressure on Sensor Accuracy	129
4.5.2	Sensor Placement and Instrumentation Limitations	130
4.5.3	Effects of Absolute Pressure and Valve Lift on Mass-Flow Agreement	132
4.5.4	Temperature Prediction Errors	133
4.5.5	Recommendations and Future Perspectives	134
5	Conclusion	138
	References	141
	Appendices	148
	Appendix 1. Representative Code Snippets from PyFlowBench Scripts	148
	Appendix 2. Uncertainty Analysis for Mass Flow Rate Measurement	157

Images

Image 1. Exhaust flow configuration at the GKN Aerospace facility. Air enters via the 'Inlet' (right), passes through the cylinder head mounted on the central test stand (labelled 'Flow bench'), and exits through the 'WG valve' to the 'Outlet' (left). Proprietary components are blurred to maintain confidentiality. 75

Figures

Figure 1. Projected marine sector energy mix under current trajectory and net-zero scenarios, visualised as a stacked area chart. Adapted based on *BP Energy Outlook 2024* (Bp Energy Outlook 2024, 2024). 16

Figure 2. Historical vs. projected increase in engine calibration complexity through addition of independent control parameters. Adapted based on (Atkinson, 2014). 20

Figure 3. Schematic representation of a vertical airflow test bench system. The layout shows key components from the test object through to the air source for an exhaust flow. Drawn by the author based on (Flow Performance, n.d.; Choate et al., 2007). 27

Figure 4. Plots of maximum intake valve law (H_{max}) and minimum intake valve law (H_{min}) showing flow velocity at the intake valve gap (W_{iv}), highlighting the substantially higher peak velocity at low lift (H_{min}) due to intensified pressure differentials across the intake valve gap. Adapted from (Clenci et al., 2014). 30

Figure 5. Comparison of normalised mean absolute error (nMAE) of velocity with Magnetic Resonance Velocimetry (MRV) data across several turbulence models. Adapted from (Nishad et al., 2019). 31

Figure 6. Comparing the velocity profiles for LES using WALE model and MRV at different Y-planes (L1-L4) at $Z=0$. From Nishad et al. (2019). 32

Figure 7. Transient and steady-state predictions of mass flow rate through the stack port, plotted against ball valve angle (left). Simulations were run at $\omega = 47$ and 94 deg/s in both clockwise (+) and counter-clockwise (-) directions. The bar chart (right) shows the averaged error (%) between transient and steady-state flow predictions. From Jeong et al. (2024). 33

- Figure 8.** Adapted plots comparing experimental and CFD-predicted mass flow rate across a range of pressure drops (ΔP), under cavitating and non-cavitating conditions (Valdés et al., 2014). 35
- Figure 9.** Comparison of simulated valve seat temperatures for different materials. The front valve seat temperature (T2) is lower for copper-based alloys, particularly brass, compared to cast iron. Data adapted from Hassan et al., 2021, Table 3. 36
- Figure 10.** Comparison of simulated front (T2) and back (T3) valve seat temperatures for five materials. Copper-based alloys show significantly better thermal performance than cast iron, with brass and beryllium copper producing the lowest seat temperatures. Data adapted from Hassan et al., 2021, Table 3. 37
- Figure 11.** Averaged heat transfer coefficient vs engine speed for intake and exhaust valves in four regions. Adapted from (Cerdoun et al., 2020). 38
- Figure 12.** Loss coefficient as a function of Reynolds number for a ball valve at three different opening positions. Data adapted from Moujaes & Jagan (2008). 39
- Figure 13.** The plots show ROAD III prototype engine's performance. The first four plots are intake discharge coefficients vs. normalised lift and the bottom four plots are exhaust discharge coefficients vs. normalised lift, redrawn from (Agnew, 1994). 44
- Figure 14.** Flow and discharge coefficients as a function of valve lift, measured on Ricardo-type with impulse torque meter, and FEV-type with paddle wheel. FEV data shows divergence at high lift due to flow choking. Adapted from El-Adawy et al. (2017). 45
- Figure 15.** Rig tumble deviation with valve lift for both flow benches. Adapted from (El-Adawy et al., 2017). 46
- Figure 16.** PIV vorticity magnitude maps at different valve lifts. From El-Adawy et al. (2017). 47
- Figure 17.** The non-dimensional rig tumble derived from the paddle wheel rig followed PIV measurements closely. Adapted from (El-Adawy et al., 2017). 47
- Figure 18.** Comparison of flow coefficients obtained from experimental, steady-state (port flow), and dynamic (cold flow) analyses. Redrawn from (Mohamad Shafie et al., 2017). 51

- Figure 19.** Ricardo tumble adaptors. From Theodorakakos (2024). 51
- Figure 20.** FEV tumble adaptor. From Theodorakakos (2024). 52
- Figure 21.** HWA and PIV velocity measuring. From Theodorakakos (2024). 52
- Figure 22.** Tumble ratio vs. Valve lift plot demonstrating how different rigs produce divergent results. Reproduced from (Theodorakakos, 2024). 53
- Figure 23.** Streamlines for HWA, Ricardo T-tube, Ricardo L-tube and the FEV tumble adaptor shown from left to right. From Theodorakakos (2024). 53
- Figure 24.** Turbulent flow visualisation in the W31 CH1 exhaust configuration, showing streamlines coloured by velocity magnitude. Warmer (red) regions indicate higher flow speeds, highlighting the complex merging and swirling of exhaust gas as they pass through the Exhaust Valves. Vectors depict the local flow direction. 61
- Figure 25.** A conical representation of the hierarchy of CFD methods, from Low-Resolution (i.e., reduced order) at the bottom to DNS at the top, demonstrating that lower resolution has lower accuracy and vice versa. Created by the author, inspired by Fig. 6.1 in Roelofs & Shams (2019). 62
- Figure 26.** Cartesian grids commonly used in the FDM. Redrawn from (Ferziger et al., 2020). 67
- Figure 27.** A cell-centred FVM grid. Control volumes are shown as squares with computational nodes (\circ) at their centres. The solid dots (\bullet) define the boundary faces where boundary conditions are applied. Redrawn from Ferziger et al. (2020). 69
- Figure 28.** Exhaust flow setup. In this configuration, the air enters Wärtsilä's inlet pipe flowing into the bottom of the cylinder liner exiting passed the exhaust valves out through the exhaust port then through the Waste Gate Valve and finally discharged through the outlet. Proprietary components are intentionally blurred in accordance with Wärtsilä confidentiality requirements. 73
- Figure 29.** Air enters the receiver section downstream of the inlet pipe, rises through the riser pipes, crosses the intake port and valves, then descends the cylinder liner and exits through the Waste Gate Valve to the outlet stack. Proprietary components are blurred in compliance with Wärtsilä confidentiality. 73

Figure 30. Schematic diagram of a pressurised flow bench system for cylinder head testing, showing airflow from the air source through the orifice plate entering inlet pipe equipped with a Pitot-static probe. It then rises through the cylinder head and finally discharges to atmosphere via the exhaust port. Arrows indicate flow direction. Drawn by the author. 74

Figure 31. Computational mesh of the intake flow CH1 fluid domain generated with SHM (left), and a cross-sectional view of the intake valves region (right). 82

Figure 32. Directory structure of a typical OpenFOAM simulation case used in this thesis. Core folders define initial conditions (0/), physical properties (constant/), and solver settings (system/). Output data are stored in postProcessing/, time directories, and parallel subdomains (processor[n]/). 85

Figure 33. Flow chart of PyFlowBench workflow. 101

Figure 34. Lifts from minimally open to fully open, exhaust flow pressure ratio of 1.1 – CH1. 105

Figure 35. CH1 exhaust flow configuration at normalised lifts of 0.25, 0.31 and 0.37 with pressure ratio of 1.5 and varying the absolute P_{in}/P_{out} . 105

Figure 36. Initial turbulence model sensitivity analysis for intake flow CH1 at normalised lift of 1.0. The figure demonstrates that the choice of experimental reference data (Wärtsilä vs. GKN) introduces a larger variation in the calculated error than the choice of turbulence model. This highlights the dominant role of experimental uncertainty in the validation process. 107

Figure 37. Five turbulence models performance against GKN Sensors at two lifts and pressure ratios. It can be observed that accuracy improved at higher lift and pressure ratio, the $k - \varepsilon$ model performed best overall, and simulations aligned more closely with GKN data than Wärtsilä's. 109

Figure 38. Five turbulence models performance against Wärtsilä sensors at two different lifts and pressure ratios. 109

Figure 39. The extended comparison of 3 turbulence models against 2 sensors at 4 distinct lifts and 2 pressure ratios. The empty rows in Wärtsilä error heat map are cases in which this sensor is not valid. It can be concluded that GKN data aligns more closely

with CFD, higher pressure ratio improves agreement, and the standard $k - \varepsilon$ model consistently outperforms the others. 111

Figure 40. Temperature validation against the experimental temperature sensor shows higher deviations at higher pressure ratio and $k - \varepsilon$ outperforming its counterparts. 112

Figure 41. Mass flow error during a pressure ratio sweep (exp. Pressure ratio 1.1–4.0). CFD matched GKN sensors better than Wärtsilä's. Accuracy improved with increasing pressure ratio. Actual pressure ratio plateaued at ~ 2.3 due to outlet backpressure. 113

Figure 42. Temperature error across the same pressure ratio sweep. Temperature predictions worsened with higher pressure ratio, indicating higher thermal inaccuracies under high-pressure conditions. 114

Figure 43. Valve lift sweep from minimally open to 100% open at pressure ratio of 1.5 with P_{in} 1.5 bar and P_{out} set to ambient pressure. The main takeaway is that with increasing lift, which corresponds to higher m , there is better matching between experimental and numerical values. 115

Figure 44. Valve lift sweep from minimally open to 100% open at pressure ratio of 1.5 with inlet pressure (P_{in}) 1.5 and outlet pressure (P_{out}) set to ambient pressure showing mass flow rate (m) normalised over lift vs their respective error percentage. Ignoring the anomalies encircled in yellow, there is a rather linear decline in the error percentage as the valve lift grows. 116

Figure 45. Valve lift sweep from minimally open to 100% open at pressure ratio of 1.5 with P_{in} 8.0 and P_{out} 5.3 bar. At higher absolute inlet pressure, there is lower uncertainty with the sensors, therefore, the deviation between the CFD and experimental values are small. 117

Figure 46. Valve lift sweep from minimally open to 100% open at pressure ratio of 1.5 with P_{in} 8.0 and P_{out} 5.3 bar showing mass flow rate (m) normalised over the lift vs its respective error percentage. The error percentage oscillates and there is no clear decrease w.r.t lift. 118

Figure 47. Temperature validation for the valve sweep at pressure ratio of 1.5 with P_{in} 1.5 bar and P_{out} set to ambient pressure. Lower absolute inlet pressure results in smaller temperature errors. 119

Figure 48. Temperature validation for the valve sweep at pressure ratio of 1.5 with P_{in} 8.0 bar and P_{out} 5.3 bar. Higher absolute inlet pressure leads to increased temperature deviation, especially at mid-range lifts. 120

Figure 49. Error percentage and mass flow rate of exhaust flow simulations at pressure ratio 1.5 with 3 turbulence models at 3 representative lifts measured against GKN's sensor. While all the turbulence models showed acceptable performance, *RNG $k - \epsilon$* outperformed at lower lifts and the error percentage for *$k - \epsilon$* decreased with increasing lifts. 122

Figure 50. Error percentage and mass flow rate of exhaust flow simulations at pressure ratio 1.5 with 3 turbulence models at 3 representative lifts measured against Wartsila's sensor. Overall, this sensor underperforms compared to GKN's. 122

Figure 51. Temperature validation for exhaust flow at pressure ratio 1.5 across 3 lifts and turbulence models. With increasing lifts, increasing temperature deviation is observed. 123

Figure 52. Side by side comparison of CH3 (left) vs CH1(right). The first two are the front and the bottom two are the back view. The high velocity region around the valve seat is more dominant in CH3 compared to CH1. Also, riser-pipe patterns are broadly similar for both and intake valve-intake valve interactions lead to higher velocity inside the cylinder. 125

Figure 53. Exhaust flow comparison of CH3 (left) vs CH1 (right) at normalised lift of 0.714. Although CH1 shows increased flow velocity at the valve seat area, CH3's higher velocity is sustained longer in the exhaust port. Also, a re-circulation zone can be observed right before the bend. 126

Figure 54. Comparison of flow structures in the CH1 intake for single-valve operation at a normalised lift of 0.257. Opening Valve A (left) versus Valve B (right) produces near-perfect mirror-image flow patterns dominated by a coherent swirl, demonstrating the port's symmetrical design. 127

- Figure 55.** CH1 exhaust valve performance – Static valve C (left) against Static valve D (right). The flow shows a characteristic swirl at the valve exit and a strong, wall-impinging jet that travels upstream into the port. 128
- Figure 56.** First batch of experimental test points for intake and exhaust configurations at a pressure ratio of 1.1. Intake tests were conducted with an inlet pressure of 1.5 bar and outlet pressure of 1.4 bar, while exhaust tests used 2.0 bar inlet and 1.8 bar outlet conditions. 130
- Figure 57.** The suboptimal location of the sensor in Wärtsilä’s inlet pipe. 131
- Figure 58.** Top view of the CH1 intake flow configuration, highlighting the bend in the y-direction where the Pitot-static tube is located. This additional bend introduces a potential source of experimental uncertainty. Proprietary components are intentionally blurred in accordance with Wärtsilä confidentiality requirements. 132
- Figure 59.** Top view of exhaust flow configuration for CH1. Wärtsilä’s inlet pipe goes straight down under the bench. Proprietary components are intentionally blurred in accordance with Wärtsilä confidentiality requirements. 132
- Figure 60.** Choosing isothermal boundary condition shows improvement in thermal prediction. The simulated temperature has been raised by 3.91 K compared to adiabatic boundary condition. 133

Tables

- Table 1.** Simulation vs. test data for $L/D = 0.25$. Data adapted from Bhowmik & Suh., 2021, Table 3. 40
- Table 2.** Comparative table of applications of 3D CFD models in valve flow research. 41
- Table 3.** Comparative table of flow bench applications in engine development. 55
- Table 4.** Flow bench test modes for full-scale W31 cylinder heads: static (0–100 % lift sweep), scavenging (\approx 5–20 % lift with both valves open), and dynamic (realistic moving lift profiles). 71
- Table 5.** Intake flow bench test matrix. 76
- Table 6.** Exhaust flow bench test matrix. 77

Table 7. Mesh quality metrics for the CH1 intake and exhaust configurations at representative low (0.14) and full (1.0) valve lifts, normalised by the maximum lift (L/L_max).	83
Table 8. Mesh-quality scorecard with key take-aways.	83
Table 9. Numerical Flow bench boundary condition reference table	86
Table 10. Main dictionaries in constant/	90
Table 11. Core dictionaries in system/ directory used in flow bench simulations conducted in OpenFOAM.	92
Table 12. A summary of baseline operating conditions.	158
Table 13. Summary of estimated uncertainties.	158
Table 14. Perturbation results.	159
Table 15. Final sensitivity analysis table.	160

Abbreviations

Acronyms

3D3C	three-dimensional, three-component
3BTDC	Before Top Dead Centre
CFD	Computational Fluid Dynamics
CH1	Cylinder Head 1
CH2	Cylinder Head 2
CH3	Cylinder Head 3
CHs	Cylinder Heads
CHT	Conjugate Heat Transfer
CVs	Control Volumes
DES	Detached-Eddy Simulations
DNS	Direct Numerical Simulation
DoE	Design-of-Experiments
FDM	Finite Difference Method
FEM	Finite Element Method
FVM	Finite Volume Method
H2ICEs	Hydrogen in Internal Combustion Engines
HCCI	Homogeneous Charge Compression Ignition
HPC	High-Performance Computing
HWA	Hot-Wire Anemometry
ICEs	Internal combustion engines
LES	Large-Eddy Simulation

ML	Machine Learning
MRV	Magnetic Resonance Velocimetry
nMAE	normalised Mean Absolute Error
PBiCGStab	Preconditioned Bi-conjugate Gradient Stabilised
PDE	Partial Differential Equation
PISO	Pressure-Implicit with Splitting of Operators
SIMPLE	Semi-Implicit Method for Pressure Linked Equations
PIV	Particle Image Velocimetry
PTV	particle Tracking Velocimetry
RANS	Reynolds-Averaged Navier-Stokes
RCCI	Reactivity Controlled Compression Ignition
RQ	Research Question
RSS	Root-Sum-Square
SHM	SnappyHexMesh
SST	Shear-Stress-Transport
URANS	Unsteady RANS
WG	Waste Gate

Greek Symbols

β	Diameter ratio d/D for an orifice
Γ	Effective diffusivity
γ	Ratio of specific heats
ε	Turbulent kinetic-energy dissipation rate
λ	Thermal conductivity
μ	Dynamic viscosity
ν	Kinematic viscosity (μ/ρ)
ρ	Fluid density
τ	Shear-stress tensor component
φ	Scalar quantities
ω	Specific dissipation rate

Other Symbols

A	Flow-passage cross-sectional area
\vec{b}	Body force per unit mass vector
C_d	Discharge coefficient
D	Up-stream pipe diameter in orifice correlation
d	Orifice (throat) diameter
e	Specific internal energy
K	ITABAR sensor calibration factor (“K-Factor”)
k	Turbulent kinetic energy
L	Valve lift
\dot{m}	Mass flow rate
P	Pressure

ΔP	Differential pressure or pressure drop
Q	Volumetric flow rate
R	Specific gas constant
Re	Reynolds number
T	Temperature
U, \vec{v}	Velocity magnitude

1 Introduction

As global pressure mounts to reduce greenhouse gas emissions and combat climate change, the marine transportation sector, like many others, faces increasing scrutiny. Internal combustion engines (ICEs), despite ongoing decarbonisation efforts, continue to power the vast majority of marine vessels units due to their high-power density, fuel flexibility, and maturity as clearly shown in Figure 1. Achieving the ambitious targets set by international regulations concerning decarbonisation and emission reduction goals requires rapid improvements to meet the 2030 and 2050 greenhouse gas reduction targets (Curran et al., 2024).

Figure 1 contrasts two trajectories for the marine sector fuel mix. In the Current Trajectory case, oil products remain pervasive, contracting only from near-total dominance in 2022 to roughly 85 % in 2035 and still exceeding 60 % by 2050, while gas and biofuels make marginal gains. In contrast, the Net Zero scenario projects a sharper oil decline falling to below 15 % by 2050, displaced principally by hydrogen-derived fuels ($\approx 40\%$) and advanced biofuels ($\approx 30\%$).

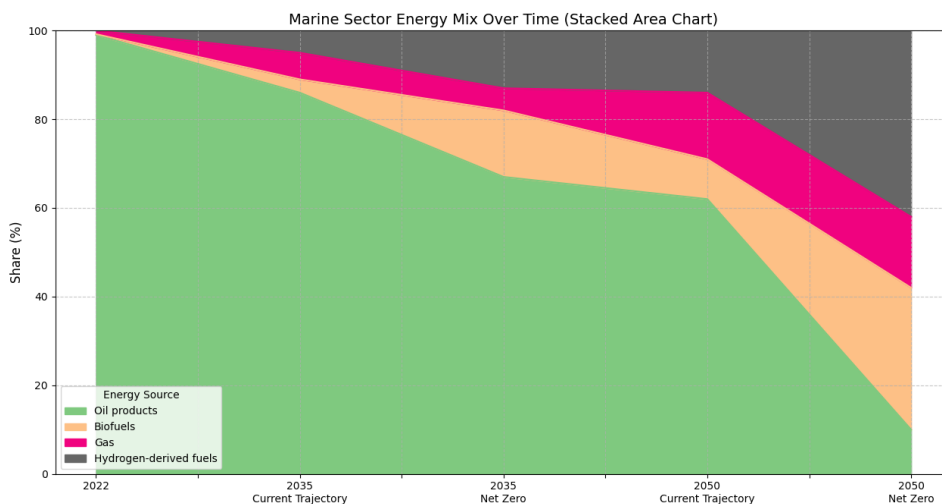


Figure 1. Projected marine sector energy mix under current trajectory and net-zero scenarios, visualised as a stacked area chart. Adapted based on *BP Energy Outlook 2024* (Bp Energy Outlook 2024, 2024).

According to Reitz et al. (2020), further improvements in ICEs can be achieved by integrating advanced combustion techniques, enhanced gas exchange, and targeted electrification, lubrication, thermal management, and after-treatment innovations. It is because irreversible processes during combustion result in a loss of around 20% of the fuel's capability to produce work. Real engines also experience heat, exhaust energy, and friction losses which range from 5-15%, 12-20% and 4-8% loss of exergy respectively (Wang et al., 2021). Therefore, improvements in these three areas can lead to considerable efficiency.

Among the most critical subsystems influencing ICE performance is the valvetrain. Intake and exhaust port geometries, valve seat angles, and timing events significantly impact cylinder filling and scavenging efficiency, which in turn dictate the combustion quality and emissions footprint (Lou & Zhu, 2020). One way to assess and improve the gas exchange capabilities of engines is through air flow bench. These test rigs allow engineers to quantify how airflow changes in response to valve lift, pressure differentials, and port designs (Eckl et al., 2008).

Hence, in order to investigate the current gas exchange potential and piston pumping work of the Wärtsilä 31 (W31) engine's cylinder heads (CHs), an experiment was conducted. To do this, a dedicated flow bench was built and transported to GKN Aerospace in Trollhättan, Sweden, with three primary objectives:

1. Product Improvement: Identify design modifications that could enhance the W31's gas exchange efficiency.
2. Method Development: Establish a robust experimental methodology for systematically evaluating different cylinder head configurations.
3. Validation Data for Simulations: Generate high-quality test data to compare against and refine computational models (Smulter, 2024).

This thesis focuses on the third point as one of its objectives, validation and verification of Computational Fluid Dynamics (CFD) simulations against experimental results.

The critical role of simulation and modelling in modern engine development stems from the unprecedented complexities engine developers now face. These include the integration of novel fuels (e.g., hydrogen, methanol, and ammonia) and advanced combustion concepts like Homogeneous Charge Compression Ignition (HCCI) and Reactivity Controlled Compression Ignition (RCCI). These specific innovations, coupled with escalating mechanical/electronic intricacy and increasingly stringent emissions regulations, create formidable multidimensional design challenges. One of such challenges is the need to re-engineer the geometry of static components since adapting to new fuels may require design modifications in cylinder head geometry to ensure air flows remain well suited to the intended fuel and combustion strategy. This thesis supports that task by assessing such flows with a numerical Flow bench.

Illustrating the diversity of new fuels, hydrogen is emerging as a promising energy carrier, primarily due to its clean combustion or conversion in fuel cells. A key area of investigation involves Hydrogen in Internal Combustion Engines (H₂ICEs), adapting traditional engines to burn hydrogen through technological advancements like manifold induction, Direct Injection (DI-H₂ICE), pressure-boosting, liquid hydrogen fuelling (L-H₂ICE), and gasoline supplementation (Singh et al., 2024). Methanol is also under active investigation; Kiouranakis et al. (2025) highlight in their review of methanol-fuelled heavy-duty engines that dual-fuel premixed strategies (e.g., diesel pilot with methanol) can significantly reduce NO_x emissions and facilitate retrofitting. Similarly, ammonia is gaining traction as a carbon-free fuel and hydrogen carrier. Its high-octane number, easier storage compared to hydrogen, and existing infrastructure make it a candidate for ICEs, particularly in sectors like marine transport. However, its direct application faces hurdles such as difficult ignition, slow flame speed, and problematic NO_x/unburned NH₃ emissions. Current research focuses on overcoming these challenges via oxygen-enriched combustion, ammonia-hydrogen mixtures (often from in-situ cracking), and plasma-assisted combustion to enhance its engine viability (Ma et al., 2023).

Alongside new fuels, advanced combustion concepts are also central to development. HCCI, for example, promises diesel-like efficiency with drastically lower NO_x and soot by igniting a lean, premixed charge via compression (Noh & No, 2017). Building on this, RCCI, a dual-fuel HCCI variant, utilises in-cylinder blending of low- and high-reactivity fuels (e.g., gasoline + diesel, or potentially ammonia + diesel) to precisely control ignition. It has emerged as one of the most efficient strategies, with recent reviews reporting RCCI achieving thermal efficiencies up to ~60% in research settings, far exceeding conventional engines, while simultaneously minimizing NO_x and soot to near-zero levels (Gowthama Krishnan et al., 2024). Therefore, novel simulation and modelling approaches are necessary to enable the efficient calibration required for managing the increased design and control requirements for continued progress in internal combustion engines (Atkinson, 2014;Reitz et al., 2020).

Figure 2 captures the escalation of calibration dimensionality by showing how many independent control (orthogonal) variables have been added to production engines over time and are expected to be added in the next decade. Historically, the industry moved from a single degree of freedom, i.e., simple injection-timing adjustment in the early 1990s, to six orthogonal variables by 2010 as exhaust-after-treatment systems came online. The projection for 2030 foresees 22 – 25 distinct parameters once advanced combustion modes, waste-heat recovery, hybridisation, full valvetrain variability, and even cycle-by-cycle control are introduced. The main takeaway from the figure is that the engine calibration dimensionality has exploded from tuning a single parameter in the 1990s to coordinating more than twenty by 2030, turning simulation-driven optimisation from a helpful tool into an operational necessity.

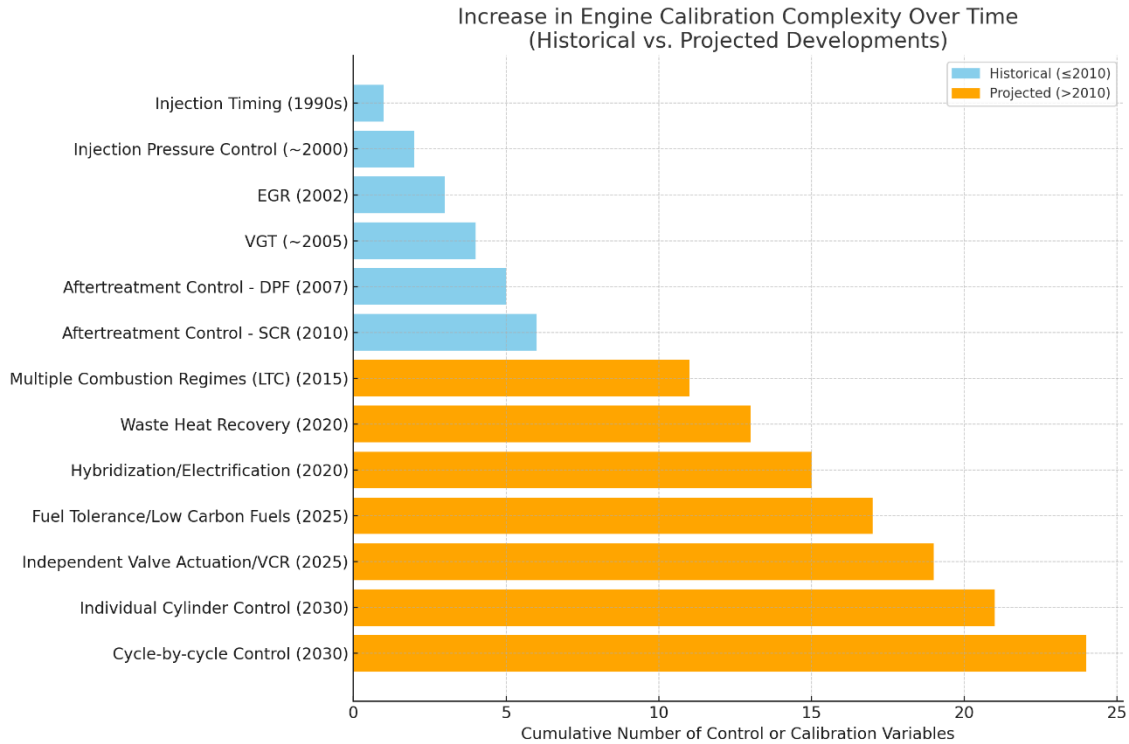


Figure 2. Historical vs. projected increase in engine calibration complexity through addition of independent control parameters. Adapted based on (Atkinson, 2014).

In addition to engine complexity, physical prototyping can be prohibitively expensive and time-consuming, therefore, computational tools have emerged as a cornerstone of modern engine development. Modelling and simulation, particularly CFD, plays a critical role in exploring engine behaviours across operational boundaries, supporting innovation while reducing development costs (Ferziger et al., 2020).

Three-dimensional CFD modelling encompasses a hierarchy of approaches for resolving turbulence, creating a distinct trade-off between fidelity and computational cost as it is further discussed in Section 2.4 and depicted in Figure 25. These approaches described by Roelofs & Shams (2019), range from fully resolved to fully modelled. At the highest level of fidelity is Direct Numerical Simulation (DNS), a method that resolves the entire range of spatial and temporal scales of turbulence without the use of any turbulence model. While DNS is considered a "numerical experiment" due to its high accuracy, it is limited by its extreme computational cost. To overcome the limitations of DNS, Large-

Eddy Simulation (LES) was developed. LES reduces the range of resolved scales by resolving only the large-scale eddies while the influence of the smaller, more universal eddies is modelled. At the most economical end of the spectrum are Reynolds-Averaged Navier-Stokes (RANS) methods, which employ time-averaging to provide a statistical description of the flow. Instead of resolving turbulent fluctuations, RANS solves for mean flow quantities only, making its computational cost only weakly dependent on the Reynolds number. Although RANS-based simulations are not nearly as computationally intensive or detailed as DNS, they are still computationally expensive, often requiring multiple hours to days of runtime per configuration depending on the availability of resources. Furthermore, setting up these simulations demands manual effort and expertise.

1.1 Thesis Goals and Research Questions

To address the challenges stated in the introduction in the context of Engine development processes at Wartsila OY, this thesis proposes an innovative workflow referred to as numerical flow bench.

The Research Goal is “To develop and validate a robust, automated CFD workflow for flow-bench simulations.” The goal is underpinned by a set of 4 research questions (RQs) that excel beyond the current engine development paradigm.

Research Question 1: How effective is workflow automation in enhancing the efficiency, productivity and consistency of engine-relevant CFD simulations?

Research Question 2: To what extent can numerical simulations accurately replicate experimental data across various flow bench configurations and operating conditions?

Research Question 3: What factors primarily influence the uncertainty budget in non-reactive, transient valve flow CFD simulations and their corresponding flow bench

experiments? Also, how can these sources of uncertainty be identified, quantified, and possibly mitigated?

Research Question 4: What potential do numerical simulations have for revealing underlying physical features that are challenging to capture through experimental measurements?

Addressing **RQ 1** and **RQ 2**, the automated and experimentally validated flow-bench simulations provide a robust platform for design-of-experiments (DoE). This platform shortens iteration cycles across multiple port designs, accelerating progress toward the ultimate goal, i.e., improving gas exchange in the W31 engine and thereby enhancing combustion efficiency, reducing fuel consumption, and lowering emissions.

To address **RQ 2** (validation) and **RQ 3** (uncertainty analysis), a Python-based automation toolkit, PyFlowBench, was developed and integrated into the flow bench CFD pipeline. This toolkit then allowed for the validation of simulations against experimental data across multiple cylinder-head geometries and operating conditions. The comparison quantifies predictive accuracy while systematically decomposing measurement and modelling uncertainties. In relation to **RQ 1** (workflow automation), the automated pipeline is benchmarked against the conventional manual process to evaluate reductions in turnaround time and inter-operator variability.

1.2 Thesis Structure

This thesis is organised into five main chapters, followed by the references and appendices.

- **Chapter 1 (Introduction)** provides the background, motivation, and outlines the thesis goals and research questions.
- **Chapter 2 (Background and Literature Review)** reviews the theoretical and methodological foundations relevant to flow bench studies and CFD modelling.

It begins with the principles and components of flow bench systems, outlining their role in assessing valve and port flow in engine development. Section 2.2 surveys valve flow research, and flow bench applications categorised into foundational, experimental, numerical, and hybrid studies. Sections 2.3 to 2.6 summarise key CFD concepts, including governing equations, turbulence modelling, and discretisation methods. This background establishes the context for the methods and tools employed in this thesis.

- **Chapter 3 (Methodology)** provides a detailed account of the experimental and numerical methods central to this research. It begins by outlining the experimental framework, covering the physical setup of the Wärtsilä 31 flow bench, the comprehensive Design of Experiments, and a rigorous uncertainty analysis of the measurement data. The chapter then explains the numerical setup, detailing the computational mesh generation strategy and the specific configuration of the CFD cases. Finally, it describes the development and core functionalities of the PyFlowBench Python package, the bespoke automation toolkit that underpins the execution of the entire simulation workflow.
- **Chapter 4 (Results)** presents the study's key findings, beginning with an initial validation. It then systematically investigates the intake and exhaust configurations through parametric and turbulence model sensitivity studies. 3D visualisations of exemplary cases provide deeper physical insight into the flow structures of different cylinder head designs. The chapter synthesises these results to diagnose the sources of experimental and numerical error, concluding with targeted recommendations for future work.
- **Chapter 5 (Conclusion)** summarises the primary outcomes of the thesis by addressing the four guiding research questions. The discussion provides evidence-based answers that integrate insights from the development and validation of the automated PyFlowBench workflow, its alignment with experimental data, its critical role in diagnosing experimental uncertainty, and its unique ability to interpret complex flow fields. The chapter concludes by reflecting on the broader implications of these findings and outlining promising directions for future work.

Following these chapters, the **References** section compiles the literature cited throughout the thesis. Finally, the **Appendices** provide additional materials and detailed information that support the main text.

2 Background and Literature Review

2.1 Principles of Flow Bench Measurements

An airflow bench is a device developed to quantify airflow through engine components, such as intake or exhaust ports, under regulated settings. The basic principle is to drive air through the test piece at a known pressure differential and measure the resulting flow rate. In a typical setup, an air pump (vacuum source) draws air through the component while a calibrated metering element (such as a sharp-edged orifice plate) measures the flow (Choate et al., 2007). In other words, the flow bench air source generates a difference in pressure across the test object, resulting in air flowing through both the test object and the entire flow bench. The air flow velocity is measured at a specific location in the flow bench. The observed air velocity allows for the calculation of the volume or mass of air flowing through the test object (*Flow Performance*, n.d.).

The volumetric flow rate Q through an orifice of area A is proportional to the square root of the pressure drop ΔP across it, where ρ is the fluid density and C is the flow coefficient expressed as:

$$Q = CA \sqrt{\frac{2\Delta P}{\rho}} . \quad (1)$$

For refined orifice, nozzle and venturi, the equation is:

$$Q = C_d A_{\text{throat}} \left[\frac{2\Delta p}{\rho(1 - \beta^4)} \right]^{\frac{1}{2}} . \quad (2)$$

In Equation 2, C_d is the dimensionless discharge coefficient which accounts for flow contraction and losses, A_{throat} is the area at the smallest cross-section (e.g. orifice hole or

venturi throat), and β is the Diameter ratio = d/D , where, d is the throat diameter and D is the upstream pipe diameter (White, 2011).

The discharge coefficient (C_d) is defined as the ratio of the actual mass flow rate through a passage to the theoretical maximum (isentropic) mass flow rate for a given reference geometry. It is often characterised for the valve flow area as follows (El-Adawy et al., 2017):

$$C_d = \frac{\dot{m}_{real}}{\dot{m}_{theor}}, \quad (3)$$

thus, C_d effectively measures how efficiently the geometry passes air compared to an ideal, frictionless passage. Despite the unsteady nature of an actual engine flow, discharge coefficients from steady-state flow bench tests have been shown to correlate reasonably well with dynamic engine breathing performance over typical operating speed ranges (Heywood, 1988) as cited by Soriano & Rech (2012). This correlation justifies the widespread use of flow benches in engine development, as they offer a practical and repeatable method to quantify port and valve aerodynamic characteristics under standardised conditions, without the complexity of full engine testing. Accordingly, this thesis centres on validating and comparing numerically predicted mass flow rates (\dot{m}) with their experimental counterparts, because \dot{m} is the quantity that directly determines the discharge coefficient C_d .

2.1.1 Typical components of a Flow Bench System

A typical airflow bench consists of several key structural and functional components, each serving a specific role in creating controlled airflow and measuring it accurately. Figure 3 shows the schematic representation of an air flow bench with the main components labelled accordingly. The components include the air source, metering element, settling chamber, bench surface, cylinder bore adapter, and the test piece or cylinder head.




Figure 3. Schematic representation of a vertical airflow test bench system. The layout shows key components from the test object through to the air source for an exhaust flow. Drawn by the author based on (Flow Performance, n.d.; Choate et al., 2007).

The air source is the core driver of airflow through the test piece. Typically implemented using a variable-speed blower or a vacuum motor. This mechanism creates a controllable pressure differential, either through suction or applied pressure, to initiate and sustain flow across the component under examination (Flow Performance, n.d.; Choate et al., 2007).

The metering element is used to determine the volumetric airflow. For differential-pressure flowmeters, it can be done by introducing a calibrated restriction that causes a measurable pressure drop or velocity change. This pressure differential is interpreted based on Bernoulli's principle or empirical calibration to calculate the final value (Flow Performance, n.d.; Choate et al., 2007). Flow measurement devices fall into three general categories (White, 2011):

1. Mechanical Devices: these measure actual fluid mass or volume by physically capturing and counting flow units. For instance, the following methods are recalled.
 - a. Mass-based methods: weighing tanks, tilting traps.

- b. Volume-based methods: reciprocating pistons, slotted rings, nutating discs, sliding vanes, gear/lobed impellers, bellows, sealed-drum compartments.
2. Differential-pressure flowmeters: These devices infer flow rate by measuring a pressure difference created by an element placed in the flow path.
- a. Head Loss Devices:
 - i. Bernoulli-type: orifice plates, flow nozzles, Venturi tubes
 - ii. Friction-based: capillary tubes, porous plugs (less common due to high head loss)
 - b. Insertion or velocity-area: Pitot tube, averaging Pitot, etc (*Flow and Level Measurement*, n.d.).
3. Modern Technologies: widely used alternatives based on different physical principles include:
- a. turbine, vortex, and ultrasonic meters
 - b. rotameters, Coriolis mass meters, laminar flow elements

These instruments are selected based on fluid type, required accuracy, and pressure loss tolerance. The metering devices used in the flow bench campaign examined in this thesis belong to the class of differential-pressure flowmeters.

To ensure a stable and uniform airflow before entering the test object, the system includes a settling chamber, also referred to as a plenum chamber. This enclosed volume mitigates turbulence and pulsations introduced by the blower or vacuum system and helps getting more accurate measurements. The bench surface is located on top of the settling chamber that serves as a mounting platform and structural interface. It enables precise positioning of the test specimen and includes a discharge port serving as a pathway linking the airflow from the flow bench to the test object (Flow Performance, n.d.; Choate et al., 2007).

In order to connect the test object, typically a cylinder head, to the plenum or bench top in a leak-free manner a fixture or adapter is required. The cylinder head is attached to a cylindrical liner that maintains the same diameter as the cylinder bore. For instance, to assess intake port flow, air is drawn through the intake port, past the valve and into the cylinder bore extension (El-Adawy et al., 2017; Flow Performance, n.d.).

With the fundamental principles and components of the flow bench now established, the following sections will review the relevant literature. The discussion will first explore the foundational physics of valve flow and then provide an in-depth survey of how the flow bench has been utilised as a critical tool in experimental, numerical, and integrated hybrid studies.

2.2 Literature Survey: Valve Flow Studies and Flow Bench Applications

2.2.1 Fundamental Studies on Valve Flow Research

Recent literature highlights how CFD contributes to understanding turbulent flow structures, unsteady behaviours, cavitation tendencies, heat transfer, and fluid–structure interactions in valve systems.

Turbulence Modelling in Valve Flows

In many valve applications turbulence plays a dual role, it can enhance mixing, which is beneficial for combustion, but it can also induce pressure losses and flow separation. A study by Clenci et al. (2014) demonstrated through 3D CFD simulation how reducing intake valve lift at idle speed significantly increased turbulence intensity and swirl within the cylinder. As shown in the Figure 4, the peak flow velocity at the valve gap increased from ~32 m/s at high lift to ~160 m/s at low lift, creating a strong shear region that

boosted turbulence kinetic energy. This led to better mixture homogenisation and improved idle stability, underscoring CFD's value in capturing subtle yet critical flow trends.




Figure 4. Plots of maximum intake valve law (Hmax) and minimum intake valve law (Hmin) showing flow velocity at the intake valve gap (W_{iv}), highlighting the substantially higher peak velocity at low lift (Hmin) due to intensified pressure differentials across the intake valve gap. Adapted from (Clenci et al., 2014).

However, the accuracy of such predictions is dependent on the chosen turbulence model. RANS models such as, the $k - \varepsilon$ model, remain widely used due to their computational affordability. While it provides acceptable accuracy in fully turbulent regimes, its isotropic eddy viscosity assumption struggles with flow separation and jet-like structures. To address this, researchers have successfully applied modified two-equation models, such as the *RNG* $k - \varepsilon$ or Shear-Stress-Transport (SST) $k - \omega$, to better handle the recirculation and swirl common in valve flows, as documented by Domagala & Fabis-Domagala (2023) and Valdés et al. (2014). Such modified models, $k - \omega$ SST for example, achieves this through two key features. Firstly, it operates as a hybrid model, blending the robust and accurate $k - \omega$ model in the near-wall region with the freestream-independent $k - \varepsilon$ model in the outer flow. Secondly, it introduces a limiter on the eddy viscosity, which enforces the physical principle that shear stress is proportional to turbulent kinetic energy, thereby improving the prediction of flow separation and recirculation (Menter, 1994).

However, traditional RANS models can fail to capture transient vortical structures and cyclic variability in intake flows. For higher fidelity, LES is employed to resolve unsteady turbulence around intake valves. The trade-off is clear: RANS remains computationally

affordable and is useful for rapid design iterations, whereas LES provides deeper insight into flow physics at a higher computational cost.

As demonstrated by Nishad et al. (2019), LES outperforms RANS in matching experimental Magnetic Resonance Velocimetry (MRV) data. Specifically, LES using the WALE model achieved a lower normalised mean absolute error (nMAE $\approx 3.48\%$) compared to RANS (nMAE $\approx 4.48\%$), as shown in Figure 5. Moreover, detailed comparisons of velocity profiles at multiple cross-sections within the cylinder confirm that LES better resolves flow gradients and jet structures across the valve curtain as shown in Figure 6.

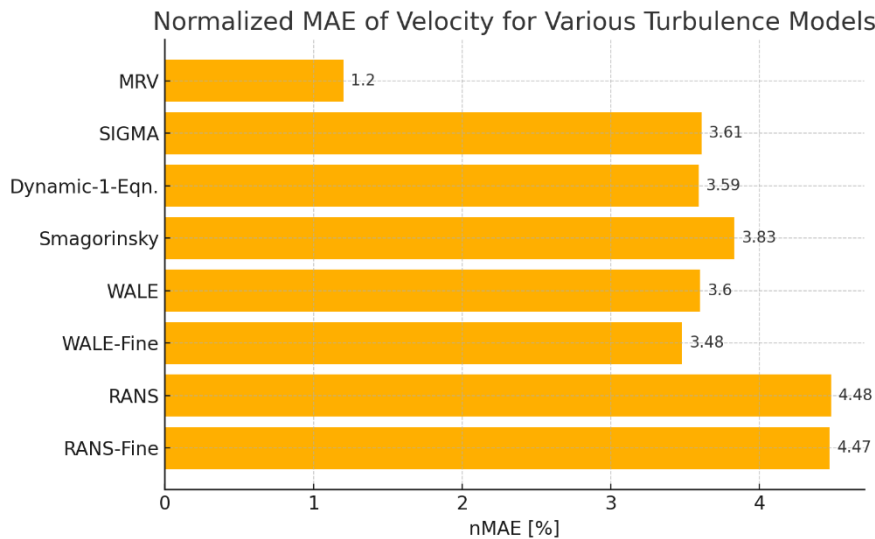


Figure 5. Comparison of normalised mean absolute error (nMAE) of velocity with Magnetic Resonance Velocimetry (MRV) data across several turbulence models. Adapted from (Nishad et al., 2019).

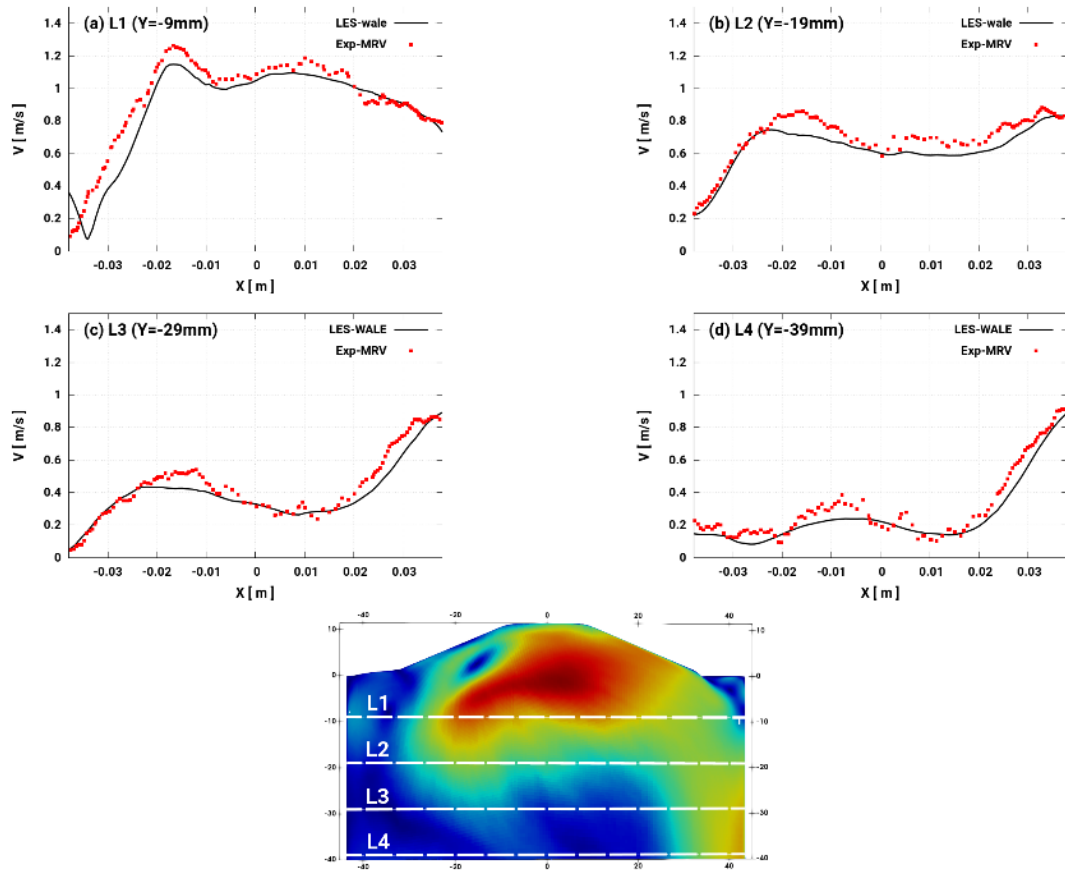


Figure 6. Comparing the velocity profiles for LES using WALE model and MRV at different Y-planes (L1-L4) at $Z=0$. From Nishad et al. (2019).

Regardless of the turbulence model used, a recurring requirement in CFD studies is robust validation against experimental data and ensuring grid independence, both of which are critical to mitigate numerical uncertainties. For instance, it was shown by Clenci et al. (2014) who carefully calibrated their intake flow CFD with boundary conditions from engine experiments resulted in good agreement in flow patterns.

Transient Flow Analysis in Valve Operation

Valves often operate under transient conditions (e.g., rapid opening/closing in engines or control systems), and CFD allows time-resolved analysis of these unsteady flows. Modern CFD solvers support moving mesh techniques that update the computational grid as the valve position changes (Domagala & Fabis-Domagala, 2023). This approach captures dynamic effects that steady-state analysis inherently misses. Although steady-

state CFD is fast and suitable for fixed-condition analyses, it cannot capture time-dependent phenomena. Whereas transient CFD resolves unsteady flow behaviour, albeit at higher computational cost, making it essential for accurately modelling dynamic systems like moving valves or pulsating flows.

In their work on a motored engine, Clenci et al. (2014) performed a dynamic simulation of the full intake stroke, using a deforming mesh to track the moving valve. Their results captured the evolution of the intake jet, backflow during valve closing, and the buildup of turbulence.

Jeong et al. (2024) provided a clear example of the disparity between transient and steady-state simulations. They performed a transient CFD study on a five-way coolant control valve using a rotating mesh to replicate the rapid actuation of the ball valve. Their analysis revealed that dynamic effects caused a flow rate deviation of up to 19.03% at the stack port when operating at a rotation speed of $\omega = 94$ deg/s, compared to steady-state predictions. As shown in Figure 7, the transient simulations captured unsteady behaviours, such as flow inertia and valve-induced hysteresis, which were not present in the steady-state results.




Figure 7. Transient and steady-state predictions of mass flow rate through the stack port, plotted against ball valve angle (left). Simulations were run at $\omega = 47$ and 94 deg/s in both clockwise (+) and counter-clockwise (-) directions. The bar chart (right) shows the averaged error (%) between transient and steady-state flow predictions. From Jeong et al. (2024).

Cavitation Prediction in Valve Flows

Cavitation, formation of vapor bubbles, can occur in regions of low pressure, such as at small valve openings or sharp throttling points. For example, Kudźma & Stosiak (2015) experimentally observed cavitation developing in a hydraulic lift valve, correlating its inception with acoustic emissions. Numerical models can visualise where vapour pockets form and collapse, helping to identify cavitation inception at geometrical singularities like seat edges or poppet gaps.

Valdés et al. (2014) validated a steady RANS CFD model of a 1.4 mm ball check valve using the Realizable $k - \varepsilon$ turbulence closure coupled with the Schnerr–Sauer homogeneous cavitation model implemented in ANSYS Fluent. Cavitation was switched on (atmospheric outlet pressure) and off (elevated outlet pressure) to mimic the test rig, and the simulation reproduced the experimentally observed $\approx 20\%$ reduction in mass flow under full cavitation as shown in Figure 8. While the Schnerr–Sauer formulation requires an empirical bubble-number density, authors therein adopted the Fluent default of 10^{13} m^{-3} and reported that the model was robust, showing less than 0.5 % sensitivity to a one-order-of-magnitude change in this value. A key insight from their validated model was the ability to calculate the hydraulic forces on the ball, showing how cavitation altered the net forces and thus the valve's stability.

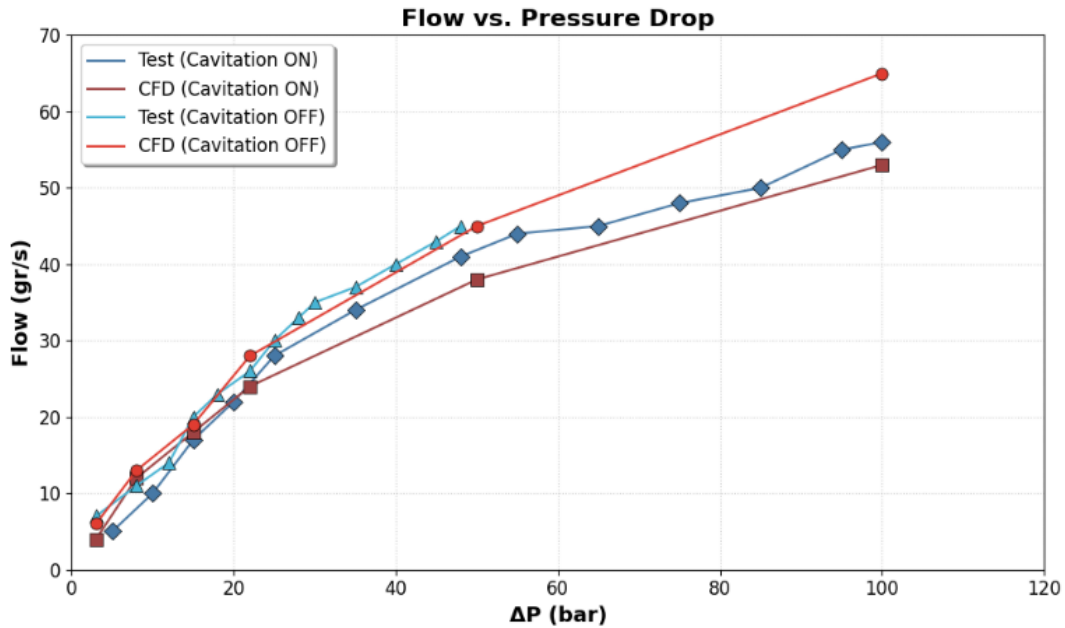


Figure 8. Adapted plots comparing experimental and CFD-predicted mass flow rate across a range of pressure drops (ΔP), under cavitating and non-cavitating conditions (Valdés et al., 2014).

Heat Transfer Assessment in Valves

Valve flows often involve significant heat transfer, especially in engine contexts where components are exposed to hot exhaust gases. Through Conjugate Heat Transfer (CHT) analysis, which couples fluid flow simulation with solid conduction models, 3D CFD can predict component temperatures by resolving detailed surface heat transfer coefficients.

For example, a transient thermal CFD study by Hassan et al. (2021) investigated the heat exchange between an engine exhaust valve and its seat to improve thermal management. Their transient thermal simulations compared a standard cast iron valve seat with four copper-based alloys. As shown in Figure 9, brass achieved a lower front seat temperature, i.e., $T_2 = 506^\circ\text{C}$ compared to cast iron, $T_2 = 538.7^\circ\text{C}$, indicating enhanced thermal conduction and improved heat removal from the valve head.

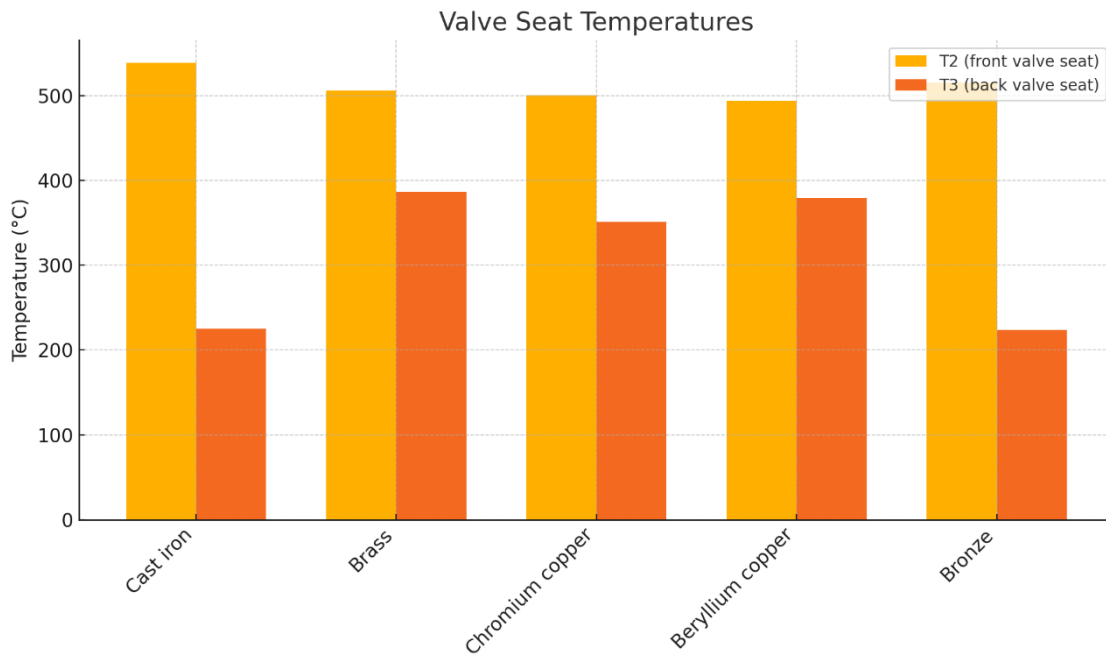


Figure 9. Comparison of simulated valve seat temperatures for different materials. The front valve seat temperature (T2) is lower for copper-based alloys, particularly brass, compared to cast iron. Data adapted from Hassan et al., 2021, Table 3.

This enhanced cooling is a direct result of the superior thermal conductivity of the copper alloys. Figure 10 quantified this effect by showing the calculated heat flux, i.e., $\frac{q_{cond}}{A} = -k \frac{dT}{dX}$, where $\frac{q_{cond}}{A}$ is the rate of heat transfer per unit area, k is the thermal conductivity of the material and $\frac{dT}{dX}$ is the temperature gradient between the valve head (T1) and the front of the seat (T2). The simulation predicted a heat flux of 14,679 kW/m² for brass, a more than three-fold increase over the 3,993 kW/m² observed for cast iron under the same conditions (Hassan et al., 2021). This supported the authors' conclusion that a brass seat can absorb on average 30% more heat in terms of temperature differences, making it a promising candidate for improved thermal management in high-temperature engine environments.

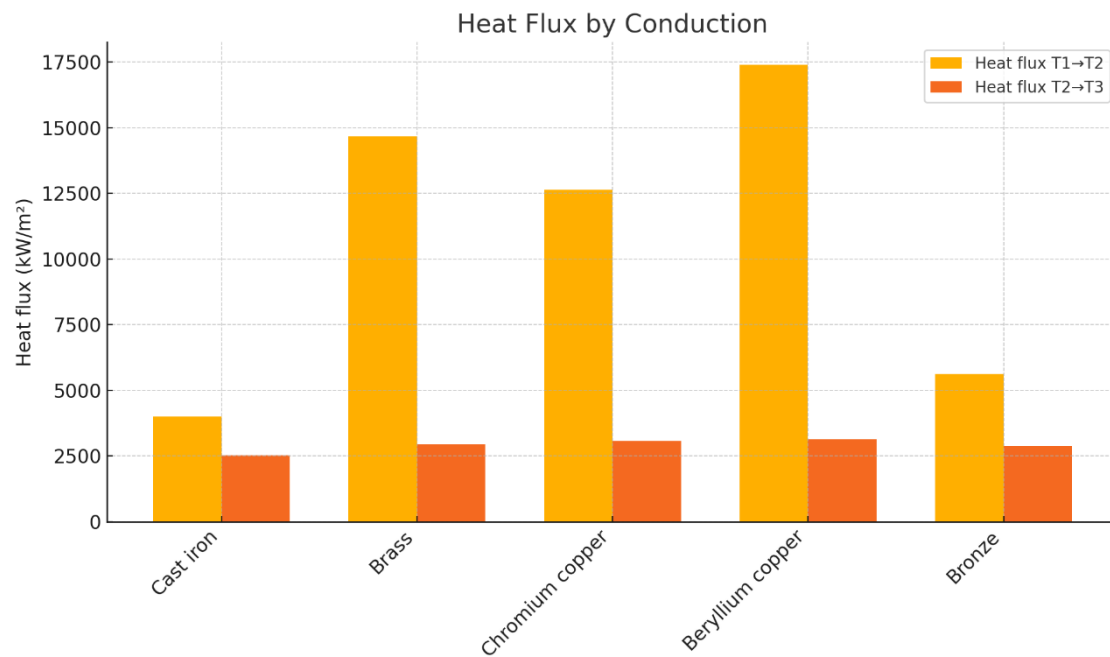


Figure 10. Comparison of simulated front (T2) and back (T3) valve seat temperatures for five materials. Copper-based alloys show significantly better thermal performance than cast iron, with brass and beryllium copper producing the lowest seat temperatures. Data adapted from Hassan et al., 2021, Table 3.

Similarly, Cerdoun et al. (2020) tackled the complex heat transfer in both intake and exhaust valves across various engine speeds. Using a coupled CFD and finite element analysis approach, they developed detailed valve temperature maps. Their results, shown in the Figure 11, demonstrated that the heat transfer coefficient varied along the valve surface, with the highest values on the seat and tulip where flow impinged or separated. A critical assumption in such CHT simulations was the accuracy of the imposed boundary conditions (e.g., fluid temperatures and convective coefficients), which can be challenging to obtain from real-world engine operation, making experimental validation crucial.

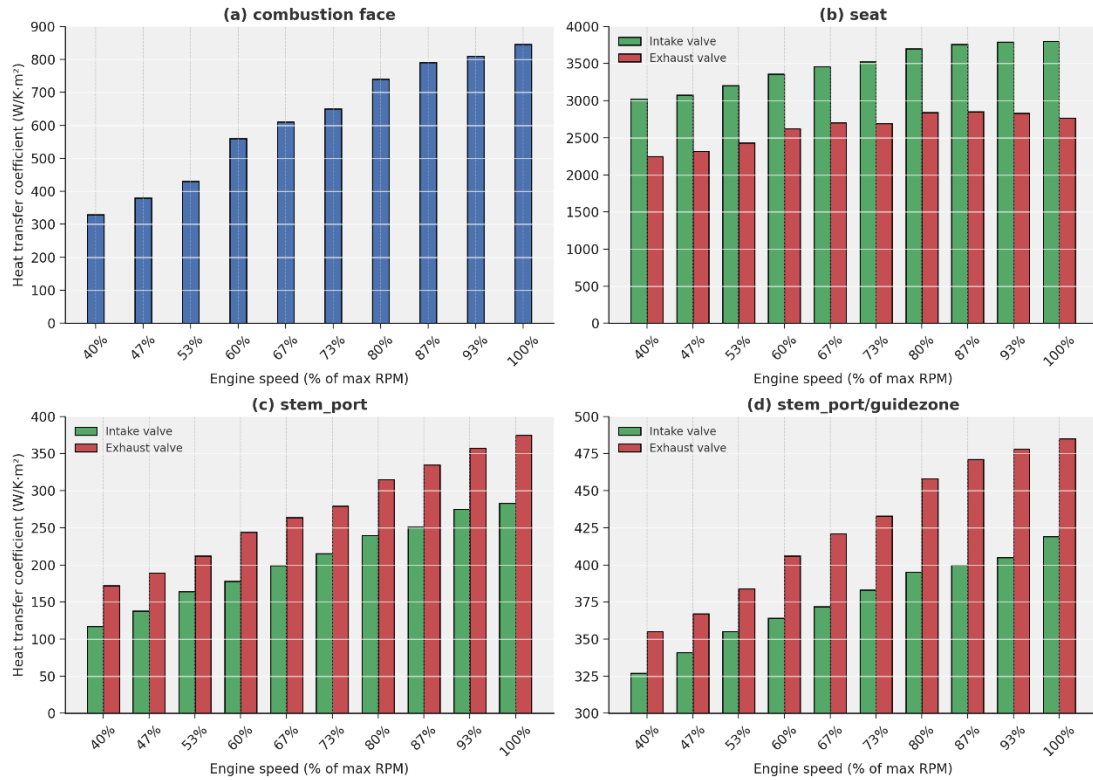


Figure 11. Averaged heat transfer coefficient vs engine speed for intake and exhaust valves in four regions. Adapted from (Cerdoun et al., 2020).

Such CFD-driven thermal analyses inform material selection and cooling strategies to prevent valve overheating. Similarly, in industrial valves, CFD-based thermal models help evaluate cooling jacket effectiveness or thermal stresses by resolving the fluid temperature fields and heat fluxes. Overall, CFD's ability to combine flow and heat transfer analysis allows engineers to assess valve thermal performance under realistic operating conditions.

General Valve Performance Prediction and Optimisation

Beyond modelling specific phenomena, 3D CFD is broadly applied to predict and optimise overall valve performance by calculating fundamental parameters, such as the loss coefficient and flow coefficient. The loss coefficient is a dimensionless measure of irreversible pressure loss and flow coefficient relates the measured mass flow rate to theoretical flow predictions. These quantities are essential for system-level design and analysis.

For instance, a study by Moujaes & Jagan (2008) used CFD to calculate these essential parameters for a ball valve under various partial openings and across a turbulent Reynolds number range of 10^5 to 10^6 . They found out that loss coefficient appeared to be largely independent of the Reynolds number for a given valve opening, as shown in the Figure 12. Their simulations, validated against experimental data, highlighted CFD's role in characterising fundamental valve hydraulic performance.

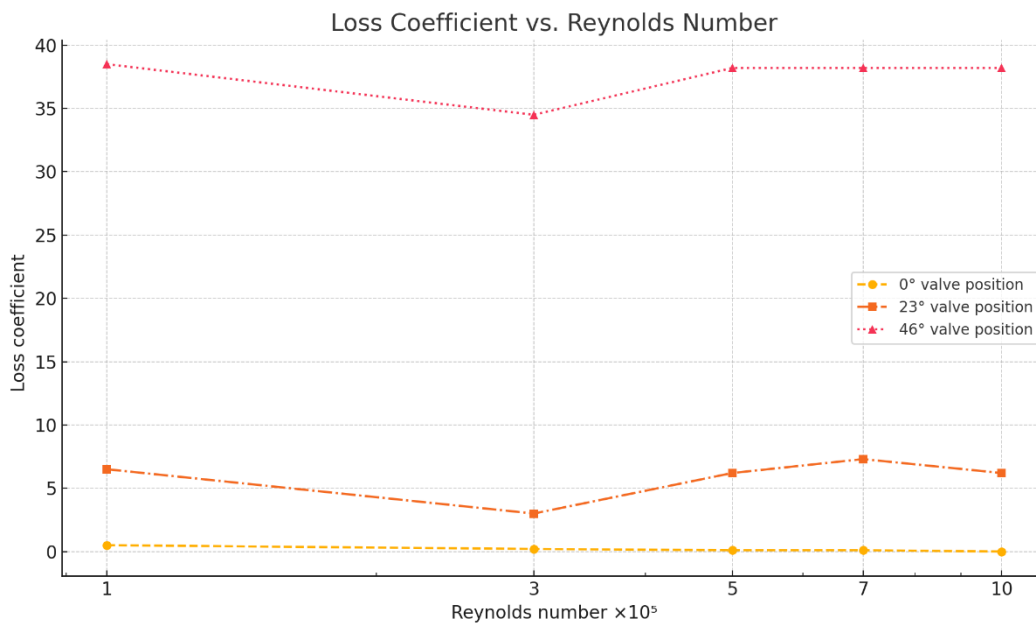


Figure 12. Loss coefficient as a function of Reynolds number for a ball valve at three different opening positions. Data adapted from Moujaes & Jagan (2008).

In another study, Bhowmik & Suh (2021) evaluated the hydrodynamic flow characteristics of a complex combined emergency stop and control valve intended for use in an ultra-supercritical steam turbine. As defined by the parameters investigated in their study, flow characteristics represent the relationship between the resulting flow rate, the valve opening (normalised as stem lift, L/D), and the system pressure ratio. They integrated theoretical analysis, full-scale 3D CFD simulations using STAR-CCM+, and experiments on a test rig.

One of the insights from this study was the challenge of experimental validation and the impact of geometric scaling. The experiments were conducted on a simplified, scaled-down version of the valve to match the test bench capabilities, while the CFD simulations were performed on the original, full-scale design. This discrepancy led to a notable scale deviation between the experimental and numerical results. As detailed in the Table 1, the simulated flow coefficient at a lift of $L/D=0.25$ was 15% higher for the control valve and 16% higher for the stop valve compared to the scaled experimental data. The authors attributed this deviation to a combination of scaling distortion, internal geometric simplifications required for the test model, and physics model simplifications (e.g., using air instead of steam).

Table 1. Simulation vs. test data for $L/D = 0.25$. Data adapted from Bhowmik & Suh., 2021, Table 3.

Scale category	Diameter (mm) – Test	Diameter (mm) – Simulate	Diameter scale ratio	Flow coefficient Test	Flow coefficient Simulate	Flow scale ratio	Scale error (%)
Control Valve	110	190	1.73	0.40	0.60	1.15	15%
Stop Vale	110	220	2.00	0.26	0.45	1.16	16%

The study highlighted a key limitation of relying on scaled-down experiments for validating the performance of complex, multi-component valves. While CFD offers the flexibility to model the true geometry, its predictions must be carefully interpreted when direct, one-to-one experimental validation is not feasible. Table 2 summarises all the articles discussed in this section.

Table 2. Comparative table of applications of 3D CFD models in valve flow research.

Reference	Geometry/Setup	Applications/Focus	Main Highlights
Clenci et al. (2014)	Motored single-cylinder spark-ignition engine, moving mesh tracks poppet intake valve at low (≈ 1 mm) vs. higher lift. Steady CFD using RANS $k - \epsilon$ turbulence closure.	Turbulence modelling in engine valve flows; impact of valve lift on in-cylinder swirl and internal exhaust-gas-recirculation	Low-lift jet increased turbulence-kinetic-energy, strengthened swirl, reduced internal exhaust-gas-recirculation and thus improved idle-mixture homogeneity and stability.
Domagala & Fabis-Domagala (2023)	Survey of industrial spool and poppet hydraulic valves; compares steady and transient CFD plus Fluid–Structure Interaction.	Unsteady flow-force mechanisms, cavitation effects, and numerical model selection.	Showed accurate prediction of valve flow forces required transient grids, SST closure and cavitation models.
Valdés et al. (2014)	1.4 mm ball-check valve; steady RANS with Schnerr–Sauer cavitation model validated against bespoke flow-bench tests.	Cavitation inception, flow-rate degradation, hydraulic-force prediction	CFD reproduced mass-flow within ± 2 % for most pressure drops; cavitation lowered flow coefficient and alters net ball forces.
Nishad et al. (2019)	Moving-mesh LES of intake-valve jet; validated via MRV.	Cycle-resolved turbulence characterisation in engine intake flows.	LES captured fine vortical structures and valve-gap turbulence missed by RANS, yielding better agreement with velocity-field measurements.
Jeong et al. (2024)	Five-way electric coolant ball valve (fuel-cell vehicle); transient rotating-mesh CFD at different actuation speeds.	Transient flow-rate deviation and thermal mixing during rapid valve actuation.	Dynamic analysis revealed up to 19 % flow-rate deviation versus steady model; flow inertia and mixing depended on ball rotation speed/direction.

Reference	Geometry/Setup	Applications/Focus	Main Highlights
Kudźma & Stosiak (2015)	Hydraulic lift poppet valve; experimental flow-visualisation and acoustics measurements	Cavitation onset and noise correlation in hydraulic valves.	Determined critical Reynolds number for cavitation; linked cavitation inception to acoustic signature, proposed real-time monitoring method.
Hassan et al. (2021)	Transient CHT analysis of exhaust valve and seat; compared cast-iron seat with four copper-alloy alternatives.	Heat-transfer assessment and material influence on valve cooling.	Brass-alloy seat absorbed $\approx 30\%$ more heat than cast iron, significantly lowering valve temperature and identifying a low-cost cooling upgrade.
Cerdoun et al. (2020)	Intake and exhaust poppet valves; boundary conditions derived from 1D gas-dynamics simulation and empirical correlations, then applied to a 2D FEM model for solid conduction across an engine-speed sweep.	Spatial Heat-Transfer-Coefficient distribution and thermal mapping in engine valves.	Seat and tulip zones exhibited peak Heat-Transfer-Coefficient; exhaust valve reached higher temperature than intake, mapped hotspots for fatigue-life assessment.
Moujaes & Jagan (2008)	5 cm flanged quarter-turn ball valve at 0° , 23° , 46° openings; steady RANS $k - \epsilon$ over Reynolds numbers $1 \times 10^5 - 1 \times 10^6$.	Valve performance prediction: loss coefficient and flow coefficient.	Showed the loss coefficient was nearly Reynolds-independent for a given opening, and flow coefficient rose steeply with opening, baseline data for pipeline models.
Bhowmik & Suh (2021)	Full-scale combined emergency-stop/control steam valve for ultra-super-critical turbine; STAR-CCM+ CFD validated on VELO hydrodynamic test rig.	Hydrodynamic flow-characteristic curves and scale-effect evaluation.	CFD reproduced flow curves within 16% of tests; showed importance of full-scale validation when scaling complex valve internals.

2.2.2 Flow Bench in Experimental Studies

Agnew (1994) explained that the diagnostic value of data generated by flow benches increase when raw volume-flow data are converted into effective area (A_e) and then normalised by key reference areas that govern the flow path: valve curtain area $A_c = \pi D_0 L$, valve head area $A_v = \pi/4 D_0^2$, and cylinder bore area $A_b = \pi/4 B^2$, where: D_0 is valve outlet head diameter and B is cylinder bore diameter. The resulting dimensionless discharge coefficients are:

$$C_{dca} = \frac{A_e}{A_c}, \quad (4)$$

$$C_{dva} = \frac{A_e}{A_v}, \quad (5)$$

$$C_{dba} = \frac{A_e}{A_b}. \quad (6)$$

These coefficients separate flow capacity from flow efficiency across the valve-lift range. As illustrated in Figure 13, Agnew benchmarked the “ROAD III” prototype against a database of 18 high-performance racing engines. This set revealed class-leading C_{dba} but poor C_{dca} and C_{dva} on the intake side, indicating oversized, shrouded inlet valves; a modest valve-diameter reduction was recommended to lift both efficiency curves without sacrificing capacity.

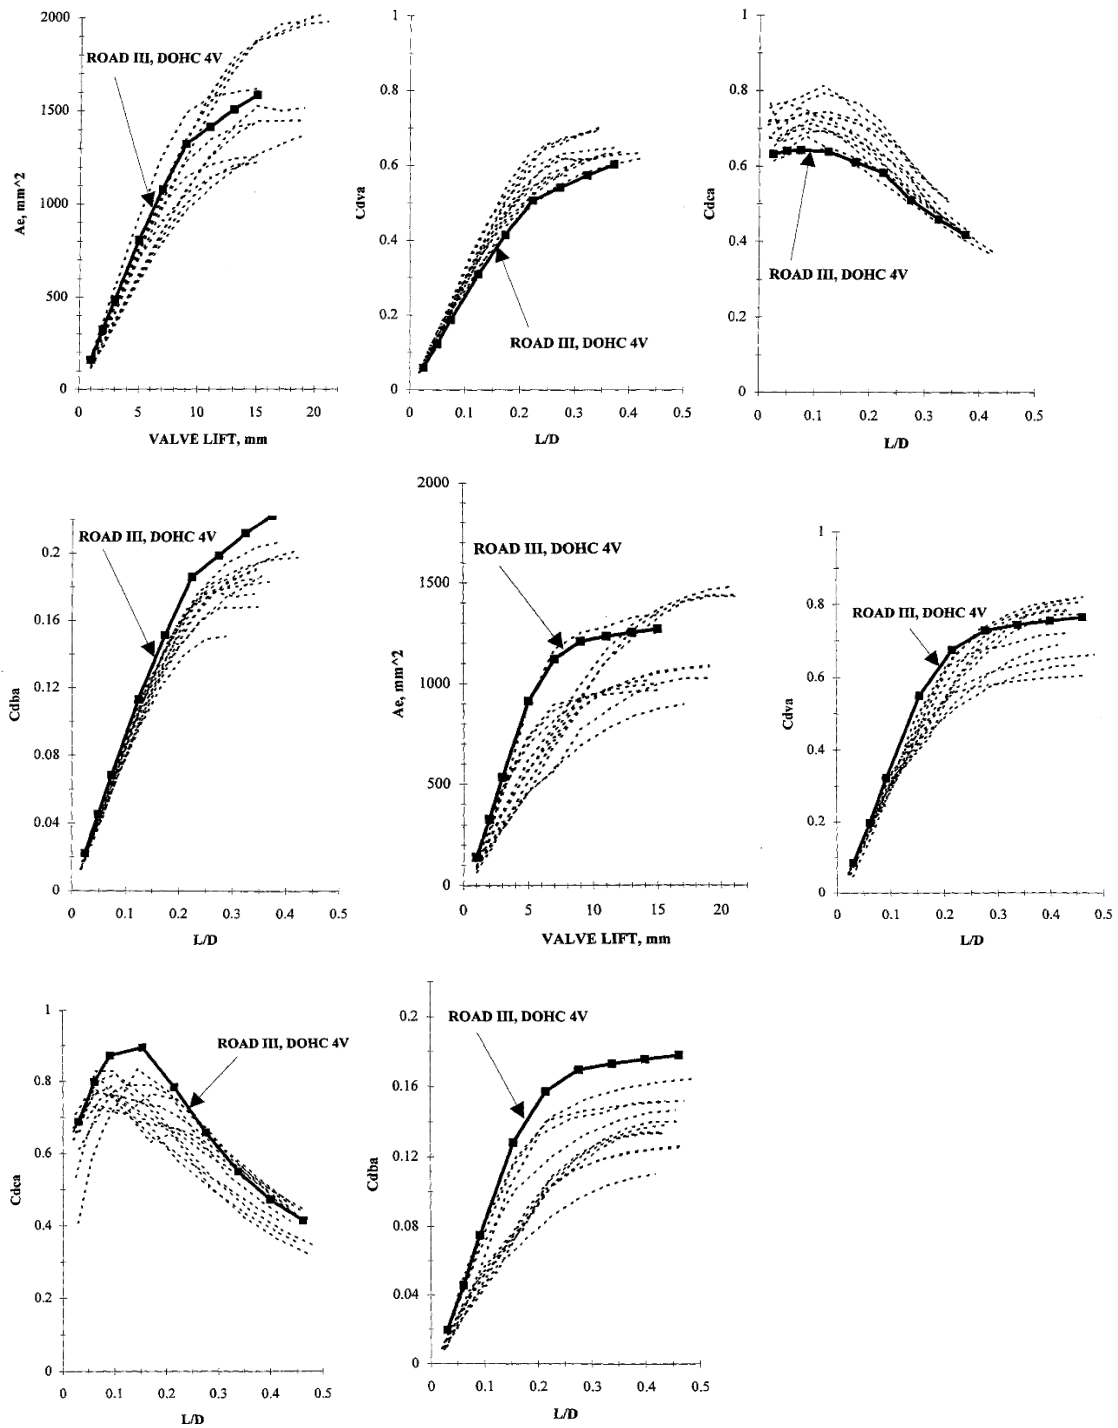


Figure 13. The plots show ROAD III prototype engine's performance. The first four plots are intake discharge coefficients vs. normalised lift and the bottom four plots are exhaust discharge coefficients vs. normalised lift, redrawn from (Agnew, 1994).

Wahono et al. (2019) identified the steady-state flow bench as a cost-effective tool for approximating bulk in-cylinder motion. They used a flow bench to investigate the intake port and valve assembly of a small (125 cm^3) single-cylinder, four-valve, four-stroke gasoline engine under various pressure differences (100 to 300 mmH₂O) and valve lifts. Key observations included an increase in air volume flow rate with both lift and pressure difference (Δp). While the flow coefficient also increased with lift, it was found to decrease with pressure difference. Conversely, the discharge coefficient was found to decrease with increasing valve lift, and showed sensitivity to pressure differences. Furthermore, a CFD model using the RNG $k - \varepsilon$ turbulence model demonstrated the progression of tumble motion inside the cylinder, underscoring the synergy between experimental flow bench data and computational modelling.

In a comparative study, El-Adawy et al. (2017) used the same four-valve cylinder head to test tumble motion in two different steady-state flow bench designs. One used an impulse torque meter (Ricardo type), and the other used a paddle wheel (FEV type). At a constant pressure difference of 600 mm H₂O they varied valve lift from 1 to 9 mm and recorded the air-flow rate, flow coefficient, discharge coefficient and non-dimensional rig-tumble. The two benches produced similar flow and discharge coefficient curves up to $\approx 6 \text{ mm}$ lift; above that point the FEV rig showed flow choking, which reduced both coefficients and altered the tumble trend as it can be observed in Figures 14 and 15.




Figure 14. Flow and discharge coefficients as a function of valve lift, measured on Ricardo-type with impulse torque meter, and FEV-type with paddle wheel. FEV data shows divergence at high lift due to flow choking. Adapted from El-Adawy et al. (2017).

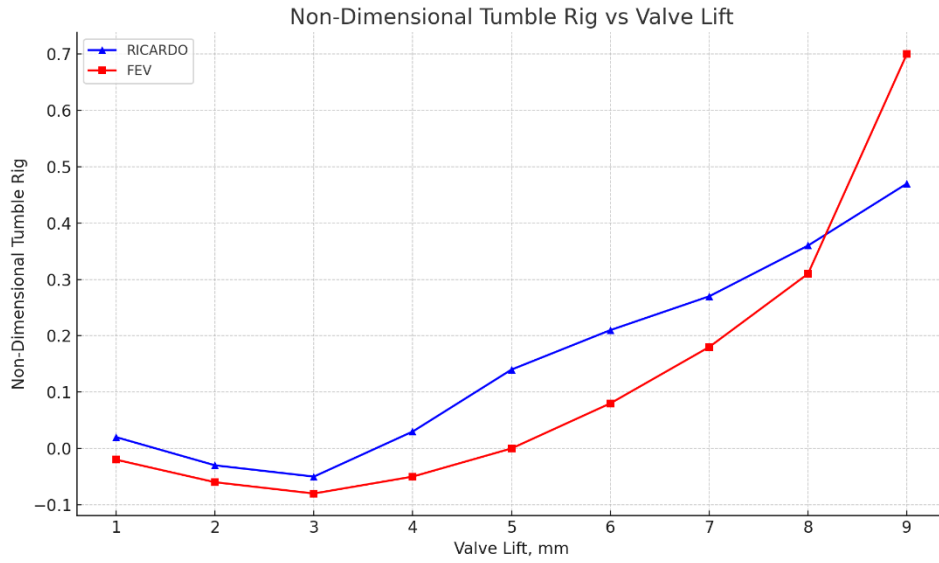


Figure 15. Rig tumble deviation with valve lift for both flow benches. Adapted from (El-Adawy et al., 2017).

They also modified the FEV-type bench to incorporate Particle Image Velocimetry (PIV), successfully visualising tumble structures and demonstrating a reasonable agreement between PIV-derived tumble and the paddle wheel measurements as quantified in Figure 17. The vorticity magnitude maps derived from PIV data are shown in Figure 16.

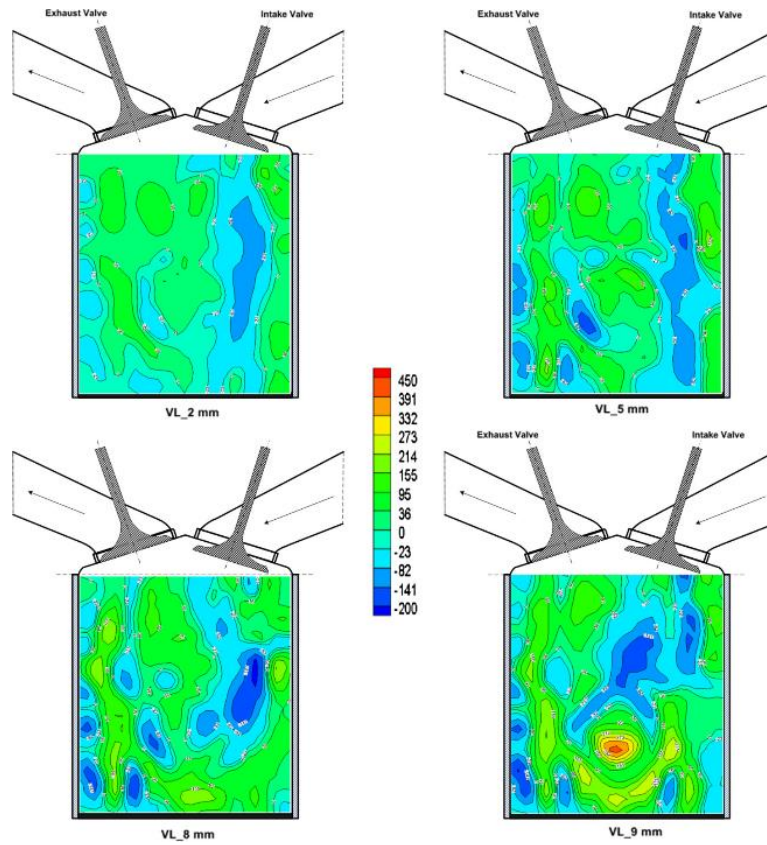


Figure 16. PIV vorticity magnitude maps at different valve lifts. From El-Adawy et al. (2017).

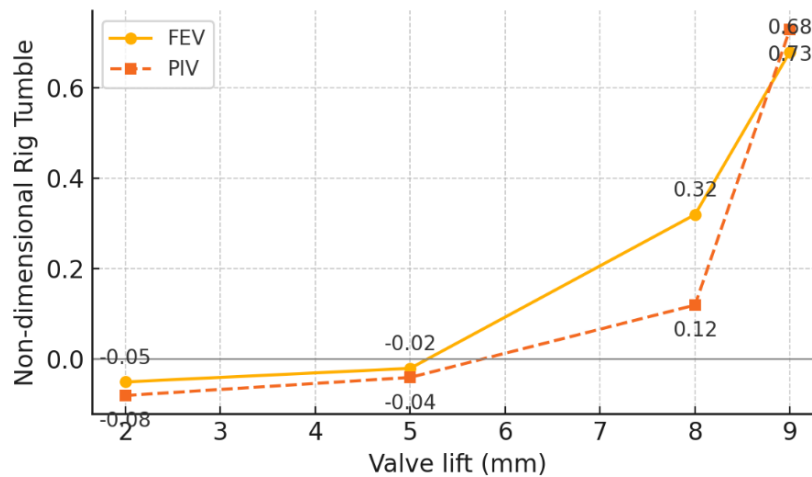


Figure 17. The non-dimensional rig tumble derived from the paddle wheel rig followed PIV measurements closely. Adapted from (El-Adawy et al., 2017).

Welch et al. (2023) also employed PIV to investigate flow behaviour. They utilised a custom-designed air flow bench to model and characterise the turbulent intake jet of an

optically accessible research engine. They removed the piston to create a simplified, continuous flow environment with high-speed PIV to investigate flow characteristics under both stationary and moving valve conditions. Their research compared flow bench data to baseline measurements from the actual motored engine, successfully simulating key features of the intake flow. A significant finding was the identification of a coherent flapping of the intake jet at 752.5 Hz, observed exclusively under the highest mass flow rate condition with stationary valves. This phenomenon was plausibly linked to vortex shedding around the valve stem, likely induced by valve vibrations. The controlled flow bench environment facilitated a comprehensive analysis of turbulence and flow instabilities. The resultant data provide robust validation for the development of CFD models.

Building on this, recent studies have introduced even more sophisticated diagnostic techniques and dynamic valve operation. Addressing the limitations of 2D measurements, Chen & Sick (2017) introduced plenoptic particle Tracking Velocimetry (PTV) on a steady-state flow bench. Plenoptic imaging captures the complete light field by recording both the position and the direction of every light ray with a microlens array, which makes depth reconstruction possible from a single view (Fang, 2025). This approach enabled three-dimensional, three-component (3D3C) velocity measurements, resolving all spatial directions and flow components. Crucially, plenoptic PTV is well-suited for optically constrained environments, such as engine cylinders, where conventional multi-camera setups are impractical. Their results successfully captured complex flow features like tumble and swirl, illustrating the potential for volumetric flow diagnostics in engine development.

2.2.3 Flow Bench in Numerical Studies

An early example of using CFD as a numerical flow bench is the work of Taghavi et al. (1990). In the first part of their study, they employed a modified KIVA code to evaluate the aerodynamic effects of minor geometric changes between two inlet port designs. The analysis quantified performance differences, finding that the original design, Port 1,

exhibited a greater head loss than the modified Port 2. The pressure deficit was approximately 750 Pa at 300° Before Top Dead Centre (BTDC), diminishing to 495 Pa by 180° BTDC. This discrepancy is due to the fact that in configuration 1 the inlet velocity hits the intake port wall with a greater angle. Despite this, Port 1 had less than 1% loss in volumetric efficiency compared to Port 2 at 180° BTDC during the intake stroke. Velocity field analysis confirmed that CFD was capable of resolving subtle geometric effects, particularly how flow was diverted due to wall interaction and valve proximity. These results demonstrated that CFD offered a level of diagnostic fidelity comparable to physical flow benches, with the added advantage of full 3D flow visualisation.

The second part of the study by Taghavi et al. extended this analysis to thermal performance, using the same CFD code to simulate in-cylinder combustion. The objective was to determine temperature and heat flux distributions on the cylinder head walls for comparison with experimental data. A key finding was the critical role of the turbulence model. The $k - \varepsilon$ turbulence model was found to outperform KIVA's initial subgrid scale model offering more accurate heat flux predictions. When validated against experimental data, the $k - \varepsilon$ based simulation showed temperature discrepancies of only about 10°C at the valve bridge and 20°C on the exhaust valve, compared to overestimations of up to 200°C on the exhaust valve head using the subgrid scale model. This confirmed its superior predictive accuracy for thermal analysis. The improvement is likely due to the two-equation $k - \varepsilon$ model solving transport equations for both turbulent kinetic energy (k) and its dissipation rate (ε), which allows the turbulence length scale to adapt dynamically and damp eddy viscosity near walls, providing critical capabilities for accurate heat transfer that the original subgrid scale model lacks.

Building on the earlier numerical flow bench study, Mohammadebrahim et al. (2012) also used CFD to analyse in-cylinder phenomena. They simulated steady-state tumble flows and validated the results against a physical flow bench with a dome-shaped cylinder head. The researchers compared seven different turbulence models and selected the RNG $k - \varepsilon$ model for its robustness, noting that the choice of model resulted in a

maximum difference of 9% in predicted vorticity. The simulation results for flow coefficient and tumble torque showed strong agreement with experimental data; the predicted flow coefficient fell within the 1% experimental uncertainty, and the tumble measurements were also in reasonably good agreement. Their analysis also revealed that the in-cylinder flow behaviour was highly sensitive to valve lift, transitioning from a uniform flow pattern at low lifts to a dominant, biased jet at higher lifts that promoted tumble.

Mohamad Shafie et al. (2017) also validated two CFD simulation methods against physical flow bench data in order to compare their accuracy. They used steady-state port flow and dynamic cold flow to determine an engine's intake breathing capacity. The authors first established baseline airflow by measuring discrete, fixed valve lifts on a flow bench. They then replicated those conditions in a static-geometry CFD simulation, termed Steady-State (Port Flow). This was followed by a transient CFD analysis that simulated the full intake stroke with moving valves and a piston, but without fuel or combustion, representing a more realistic engine cycle that is referred to as "Dynamic (Cold Flow)." Their key findings were that port flow simulations were highly accurate, while dynamic simulations consistently underpredicted breathing capacity. They recorded flow coefficient deviations between 1.6% and 5.1% for the port flow simulations, while the dynamic cold flow simulations deviated by as much as 27%, as illustrated in Figure 18. The authors attributed this discrepancy to the transient flow losses associated with valve and piston motion, which are not captured in static analyses. This study concluded that while steady-state simulations are superior for predicting absolute performance values, dynamic simulations are valuable for establishing performance trends and a lower-bound for breathing capacity.




Figure 18. Comparison of flow coefficients obtained from experimental, steady-state (port flow), and dynamic (cold flow) analyses. Redrawn from (Mohamad Shafie et al., 2017).

To address the uncertainties posed by different steady-state flow rig configurations, Theodorakakos (2024) used CFD to numerically simulate and compare four rigs used to measure the tumble motion in a four-valve, pent-roof engine cylinder head. The methods analysed were the Ricardo T-tube, Ricardo L-tube, FEV, and Hot-Wire Anemometry (HWA). The four widely used tumble measurement configurations shown schematically in Figures. 19, 20 and 21.




Figure 19. Ricardo tumble adaptors. From Theodorakakos (2024).




Figure 20. FEV tumble adaptor. From Theodorakakos (2024).




Figure 21. HWA and PIV velocity measuring. From Theodorakakos (2024).

Their primary finding was that the choice of rig significantly impacted the measured flow coefficient and tumble ratio because the rig's geometry fundamentally affects the in-cylinder flow structure. The study found that each configuration presented distinct advantages and limitations leading to different tumble ratio predictions across the valve lift range as shown in Figure 22. The FEV rig, with its small-diameter exit tube, restricted the mass flow rate but generated a strong, well-defined tumble vortex, making it suitable for assessing the influence of a real piston crown. On the other hand, the Ricardo T-tube and L-tube configurations showed divergent behaviour at high valve lifts. The T-tube developed a secondary, counter-rotating vortex that changed the flow dynamics, while the L-tube produced a stronger but asymmetric primary vortex. The HWA method was identified as an "indirect" technique that measures axial velocity non-uniformity rather than a

coherent tumble structure, yielding results that only align with other methods at high valve lifts. Figure 23 shows the formation of tumble in the four setups.

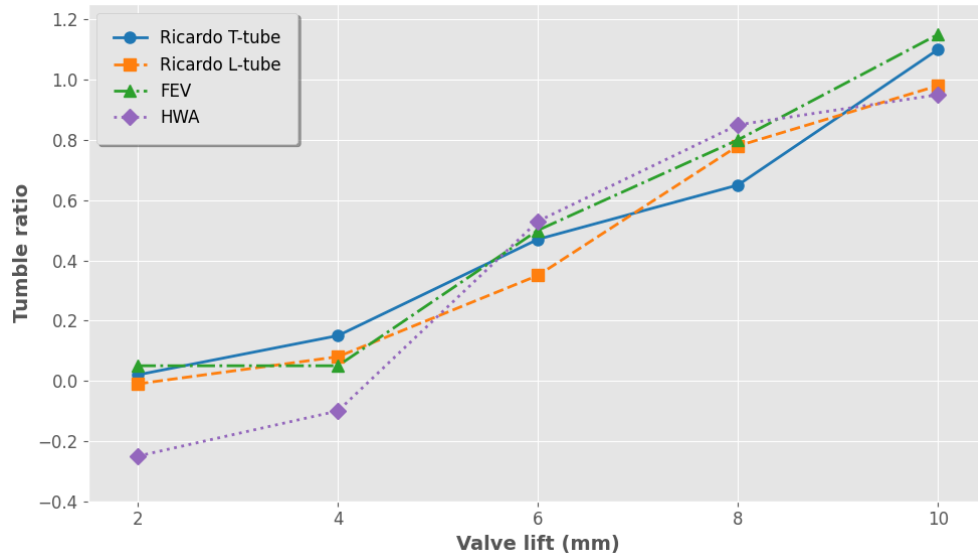


Figure 22. Tumble ratio vs. Valve lift plot demonstrating how different rigs produce divergent results. Reproduced from (Theodorakakos, 2024).

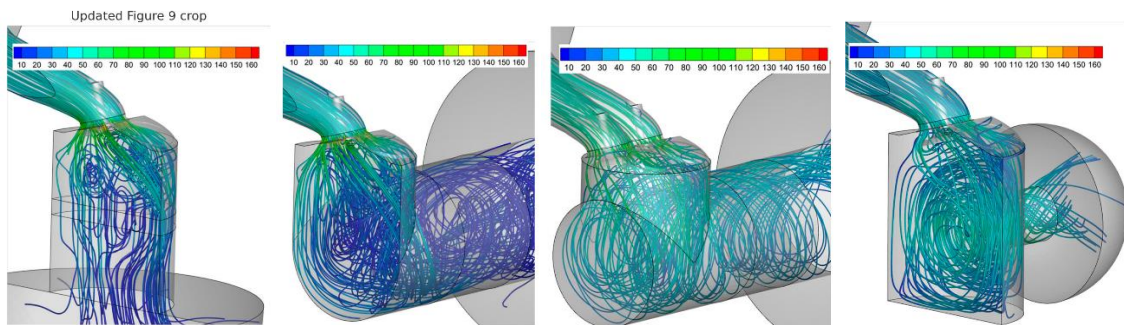


Figure 23. Streamlines for HWA, Ricardo T-tube, Ricardo L-tube and the FEV tumble adaptor shown from left to right. From Theodorakakos (2024).

2.2.4 Flow Bench in Hybrid Experimental-Numerical Studies

The trend of integrating detailed optical diagnostics with flow benches for rigorous CFD validation is further exemplified by Hartmann et al. (2016). They combined highly resolved two-dimensional, two-component PIV measurements on a steady-state flow

bench with Detached-Eddy Simulations (DES) to characterise engine-relevant flow features in the vicinity of the intake valve at a fixed lift. DES is a hybrid modelling approach that combines the strengths of two other methods: RANS and LES. Hartmann et al. (2016) work focused on comparing mean and fluctuating velocity fields, jet orientation, and critically, the validity of wall functions near the valve seat and the transport of vortex shedding from the valve stem into the combustion chamber. While general agreement was found between the finer DES mesh and PIV data, discrepancies in jet orientation were noted, possibly due to differences between the idealised CFD geometry and real-world manufacturing tolerances. The study also revealed that standard wall functions were inadequate near the valve seat in the simulations and numerically identified valve stem vortex shedding that could persist into the cylinder, highlighting the flow bench as a demanding test case for scale-resolving CFD models and their underlying assumptions.

The practical application of this integrated experimental-numerical flow bench approach to address specific engineering challenges in diesel engine development was demonstrated by Demirkesen et al. (2020). They focused on characterising swirl motion in heavy-duty diesel engine cylinders, using steady-state flow bench testing (torque meter for swirl ratio), PIV for flow visualisation, and CFD simulations (RNG $k - \varepsilon$ with adaptive mesh refinement). A key contribution of their work was the investigation of cylinder-to-cylinder variations and the significant impact of manufacturing imperfections, such as surface roughness and valve tilt, on swirl characteristics. By comparing cast cylinder heads with a smooth, rapid-prototyped flow box, and through endoscopic inspection, they quantified how these real-world factors can alter flow behaviour and deviate from idealised numerical models. Their CFD model, validated against experimental mass flow and swirl ratios, served as a tool for port design, while the experimental findings underscored the necessity of considering manufacturing quality in achieving desired in-cylinder flow. A summary of all the studies reviewed thus far has been presented in Table 3.

Table 3. Comparative table of flow bench applications in engine development.

Study	Geometry / Setup	Application / Purpose	Methods / Tools Used	Main Highlights / Findings
Taghavi et al. (1990)	CFD of intake port designs (Port 1 vs. Port 2)	Simulated physical flow bench for early port design evaluation	KIVA code, $k - \epsilon$ turbulence model	Port 1 had slightly higher head loss but <1% volumetric efficiency difference; CFD resolved 3D flow and subtle design effects
Agnew (1994)	Multiple race engine cylinder heads on steady flow bench	Normalised and benchmarked intake efficiency across engines	Steady flow bench, A_e , C_{dca} , C_{dva} , C_{dba} , L/D plots	Introduced normalised flow coefficients; built comparative database to diagnose design flaws
(Mohammedbrahim et al., 2012)	Dome-shaped cylinder head; port-valve-liner test rig	Validated CFD-predicted tumble flow vs. experiment	Fluent CFD (7 turbulence models), swirl meter, flow bench	Strong agreement (<1% for flow coefficient, ~9% for tumble); tumble increased at high lifts; RNG $k - \epsilon$ most robust
Hartmann et al. (2016)	Intake port of a spark-ignited single-cylinder head on flow bench; PIV plane and fixed lift	Assessed intake jet features and validated DES against PIV	DES (ANSYS), 2D-2C PIV, wall function assessment	DES matched average flow; stem vortex shedding observed; wall functions failed near valve seat
El-Adawy et al. (2017)	4-valve cylinder head tested on Ricardo & FEV benches	Compared tumble measurement techniques & incorporate PIV	Paddle wheel, torque meter, valve lift sweeps, PIV	Good agreement on discharge coef. up to ~6 mm; tumble results diverged between benches; PIV visualised vortex; piston presence influenced data
Chen & Sick (2017)	4 valve Spark-ignition-direct-ignition engine cylinder head with plenoptic camera PTV setup	Captured full 3D3C flow fields in realistic geometry	Single-camera PTV, RxFlow 3.1, laser optics	Visualised full tumble/swirl in 3D; asymmetric intake jet captured; overcame 2D-PIV limitations

Study	Geometry / Setup	Application / Purpose	Methods / Tools Used	Main Highlights / Findings
Mohamad Shafie et al. (2017)	CAMPRO 1.6-L intake; flow bench baseline + steady ("port-flow") & transient ("cold-flow") CFD	Compared accuracy of static vs dynamic CFD to flow-bench data	Flow bench (fixed lifts); Fluent steady RANS & moving-mesh transient	Port-flow CFD within 1.6–5.1 % of test; cold-flow underpredicted by ≤ 27 %; steady CFD best for absolute capacity, transient CFD good for trend/ lower-bound
Wahono et al. (2019)	125 cm ³ , 4-valve gasoline engine head	Assessed influence of lift and ΔP on flow & tumble	Steady flow bench (100–300 mmH ₂ O), CFD (RNG $k - \varepsilon$)	Flow rate increased with lift; flow coef. decreased with ΔP ; discharged coef. decreased with both; tumble growth captured via CFD
Demirkesen et al. (2020)	3 diesel heads + rapid prototype flow box	Hybrid swirl study: flow bench + endoscopy + CFD	Torque-based swirl bench, PIV, Converge CFD (adaptive mesh refinement, RNG $k - \varepsilon$)	Mass flow error <5%; swirl error <10%; valve tilt and surface roughness caused reversed swirl
Welch et al. (2023)	Optical bench; motored engine geometry without piston	Characterised turbulent intake jet and unsteady flapping	High-speed PIV, fixed valve lift, motored engine baseline	Coherent jet flapping at 752.5 Hz observed; attributed to valve stem vortex shedding; matched motored engine behaviour
Theodorakakos (2024)	4-valve pent-roof head; four tumble rigs (Ricardo T/L-tube, FEV, HWA) simulated	Quantified rig-geometry effect on flow-coefficient & tumble ratio	CONVERGE CFD; valve-lift sweep (2–10 mm)	Rig choice changes results: FEV had strong centred vortex but low mass flow; T-tube developed secondary counter-vortex; L-tube yielded stronger asymmetric vortex; HWA method aligned only at high lifts

2.3 Governing Conservation Equations

The behaviour of a fluid flow is governed by three fundamental conservation laws of physics: the conservation of mass, momentum, and energy. Together, these principles form a set of coupled, non-linear partial differential equations. The Eulerian, integral conservation statement for an arbitrary scalar φ over a control volume V reads (Montorfano, 2017)

$$\frac{\partial}{\partial t} \int_V \rho \varphi dV + \int_S \rho \varphi \vec{v} \cdot \vec{n} dS = 0. \quad (7)$$

In the above equation, ρ is the fluid density, \vec{v} is the velocity vector field, $\rho\varphi$ is interpreted as the “amount” of φ per unit volume. This equation forms the basis for the specific conservation laws summarised below.

2.3.1 Mass

The law of conservation of mass asserts that mass cannot be destroyed or created. It is also referred to as the continuity equation. Setting $\varphi = 1$ in Eq. (7) gives the continuity equation, i.e.,

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{v}) = 0. \quad (8)$$

So, the continuity equation in Cartesian coordinates is:

$$\frac{\partial \rho}{\partial t} + \frac{\partial(\rho u_i)}{\partial x_i} = 0. \quad (9)$$

For incompressible flow, $\nabla \cdot \vec{v} = 0$.

2.3.2 Momentum

The conservation of momentum is an expression of Newton's second law of motion applied to a fluid element. It states that the rate of change of momentum of a fluid particle is equal to the sum of the forces acting upon it. These forces include surface forces, such as pressure and viscous stresses, and body forces, such as gravity. The momentum of a fixed control volume is conserved according to the integral form shown in Equation (10).

$$\frac{\partial}{\partial t} \int_V \rho \vec{v} dV + \int_S \rho \vec{v} (\vec{v} \cdot \vec{n}) dS = \int_S T \cdot \vec{n} dS + \int_V \rho \vec{b} dV. \quad (10)$$

Here, T represents the stress tensor and \vec{b} represents the body force per unit mass. The corresponding differential form, often referred to as the "divergence form" used in finite-volume CFD codes, is given in Equation (11).

$$\frac{\partial(\rho \mathbf{v})}{\partial t} + \nabla \cdot (\rho \mathbf{v} \mathbf{v}) = \nabla \cdot \mathbf{T} + \rho \mathbf{b}. \quad (11)$$

This form directly enforces momentum conservation on each control volume when discretised, leading to robust global conservation. For the i –th component in a Cartesian coordinate system, Equation (11) can be written as shown in Equation (12), where τ_{ij} is the viscous stress tensor, p is pressure, and g_i is the gravitational body force component.

$$\frac{\partial(\rho u_i)}{\partial t} + \frac{\partial(\rho u_j u_i)}{\partial x_j} = \frac{\partial(\tau_{ij})}{\partial x_j} - \frac{\partial p}{\partial x_i} + \rho g_i. \quad (12)$$

2.3.3 Energy

The conservation of energy is derived from the first law of thermodynamics and accounts for the transport and transformation of internal, kinetic, and potential energy within a

fluid. In CFD, the total energy per unit mass, e_t , is typically expressed as the sum of internal energy e and kinetic energy v^2 , where \vec{v} is the fluid velocity vector.

The differential form of the total energy conservation equation, expressed in conservative form is (Anderson, 2009):

$$\frac{\partial}{\partial t} \left[\rho \left(e + \frac{v^2}{2} \right) \right] + \nabla \cdot \left[\rho \left(e + \frac{v^2}{2} \right) \vec{v} \right] = \nabla \cdot (k \nabla T) - \nabla \cdot (p \vec{v}) + \nabla \cdot (\boldsymbol{\tau} \cdot \vec{v}) + \rho \vec{f} \cdot \vec{v}, \quad (13)$$

where, ρ is the fluid density, p represents the pressure, k signifies the thermal conductivity, $\boldsymbol{\tau}$ indicates the viscous stress tensor, \vec{f} refers to the body force per unit mass, and T stands for the temperature. This formulation captures the essential contributions to total energy change from thermal conduction, pressure work, viscous dissipation, and external body forces. In this thesis, the sensible enthalpy has been considered the Energy variable in the OpenFOAM CFD code simulations.

2.3.4 Scalar Quantities

In addition to mass, momentum, and energy, many CFD problems require solving for the transport of other scalar quantities, φ , such as chemical species concentration, or turbulent kinetic energy. These scalars are governed by a general transport equation. The integral form is given in Equation (14), and the corresponding differential form is in Equation (15).

$$\frac{\partial}{\partial t} \int_V \rho \varphi dV + \int_S \rho \varphi (\vec{v} \cdot \vec{n}) dS = \int_S \Gamma \nabla \varphi \cdot \vec{n} dS + \int_V q_\varphi dV. \quad (14)$$

$$\frac{\partial(\rho \varphi)}{\partial t} + \nabla \cdot (\rho \vec{v} \varphi) = \nabla \cdot (\Gamma \nabla \varphi) + q_\varphi. \quad (15)$$

In these equations, Γ represents the effective diffusivity of the scalar φ , and q_φ is the source term per unit volume.

It is worth noting that the system of equations for momentum, energy, and scalar transport requires closure for turbulent fluctuations, the topic of which is discussed in Section 2.5.1.

2.3.5 Equation of State

To close the system of governing equations, a thermodynamic relationship between pressure, density, and temperature is required. This is provided by an equation of state. For many gases at low to moderate pressures, the ideal gas law, presented in Equation (16), is a suitable approximation.

$$p = \rho RT. \quad (16)$$

In Equation (16), p is the absolute pressure, ρ is the density, T is the absolute temperature, and R is the specific gas constant.

2.4 Overview of Turbulence Modelling

The majority of fluid flows encountered in practical engineering applications are characterised by turbulence, a flow regime exhibiting distinct features compared to the laminar flows. Such features include the flow unsteadiness, three-dimensionality, high vorticity, coherent structures, wide range of scales, enhanced mixing, and dissipation. These turbulent fluctuations enhance mixing and transport of momentum, heat, and mass, but also increase frictional drag and energy dissipation (Ferziger et al., 2020).




Figure 24. Turbulent flow visualisation in the W31 CH1 exhaust configuration, showing streamlines coloured by velocity magnitude. Warmer (red) regions indicate higher flow speeds, highlighting the complex merging and swirling of exhaust gas as they pass through the Exhaust Valves. Vectors depict the local flow direction.

A key challenge in 3D CFD is representing the effects of turbulence, which can be handled using a hierarchy of turbulence modelling approaches. As mentioned by Ferziger et al. (2020), various modelling approaches have been developed which offer a trade-off between fidelity and computational cost. The appropriate method depends on the desired accuracy, physical detail, and available computational resources.

Accordingly, such approaches include RANS models for mean-flow prediction and their closure, LES for resolving large-scale turbulent eddies while modelling the smaller ones, and DNS for full-resolution turbulence modelling. Figure 25 illustrates this hierarchy of

CFD modelling techniques, beginning with low-resolution or reduced-order methods at the base and progressing toward high-fidelity approaches such as DNS at the apex matching the accuracy with the computational expense.

In many cases, however, RANS predictions are sufficiently accurate for engineering and industrial purposes. For instance, NASA’s turbulence modelling workshop findings showed that RANS can be “robust and accurate for a wide variety of turbulent flows”, forming “the basis of the majority of engineering CFD simulations”(Bush et al., 2019). Furthermore, RANS-based 3D CFD computations “strike a balance between computational accuracy and cost-effectiveness” in industrial applications involving complex geometries (Cutrone et al., 2024). Given the accuracy–cost trade-offs highlighted above and the Flow Bench’s requirement for rapid prototyping, RANS offers the most suitable balance. Accordingly, this thesis employs a RANS-based approach.

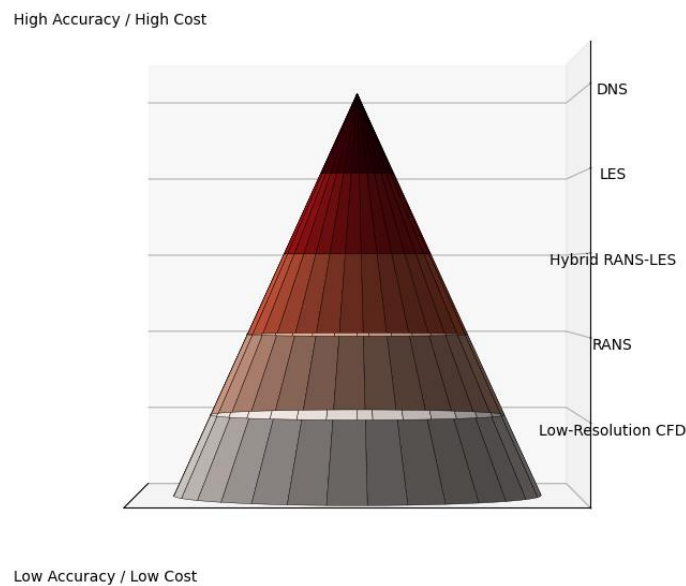


Figure 25. A conical representation of the hierarchy of CFD methods, from Low-Resolution (i.e., reduced order) at the bottom to DNS at the top, demonstrating that lower resolution has lower accuracy and vice versa. Created by the author, inspired by Fig. 6.1 in Roelofs & Shams (2019).

2.4.1 RANS Models and Equations

For many engineering applications, the primary interest lies in the time-averaged properties of a flow, such as mean forces, pressures, and heat transfer rates, rather than the instantaneous turbulent fluctuations. The RANS approach provides a computationally efficient framework for predicting these mean quantities and remains the most widely used method for industrial CFD simulations. RANS averages the governing Navier-Stokes equations in two primary ways, i.e., time averaging and ensemble averaging.

The time-averaged value is obtained by decomposing any flow variable φ into a time-averaged mean ($\bar{\varphi}$) and a fluctuation (φ'), as shown in Equation (17).

$$\varphi(x_i, t) = \bar{\varphi}(x_i) + \varphi'(x_i, t). \quad (17)$$

The mean is obtained by integrating over a time interval T long enough compared to the fluctuation timescales, where:

$$\bar{\varphi}(x_i) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \varphi(x_i, t) dt. \quad (18)$$

By definition, the average of the fluctuation is zero ($\bar{\varphi}' = 0$). This process results in steady governing equations for the mean flow, known as the RANS equations.

Alternatively, ensemble averaging is used for statistically unsteady flows, where deterministic, coherent structures might exist alongside random fluctuations. The average is taken over a large number (N) of statistically identical realizations (an ensemble) of the flow described in Equation (19),

$$\bar{\varphi}(x_i, t) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \varphi(x_i, t). \quad (19)$$

This averages out the random turbulent fluctuations while potentially preserving large-scale, repeatable unsteadiness. The resulting equations govern the unsteady mean flow and are referred to as Unsteady RANS (URANS) or Transient RANS (TRANS).

The crucial consequence of averaging is the appearance of new terms known as the Reynolds stresses ($\overline{\rho \hat{u}_i \hat{u}_j}$) in the momentum equation and the turbulent scalar fluxes ($\overline{\rho \hat{u}_i \hat{\phi}}$) in the scalar transport equation. These terms represent the transport of momentum and scalars due to turbulent fluctuations and are unknown correlations. The fact that there are more unknowns than equations is known as the "closure problem".

A common approach to modelling the Reynolds stresses and turbulent fluxes draws an analogy with molecular transport, assuming turbulence enhances viscosity and diffusivity. Boussinesq Hypothesis and Gradient Diffusion Hypothesis reduce the unknowns from the six independent Reynolds stresses and three flux components to just μ_t and Γ_t . These two can be determined by Zero-Equation Models such as Prandtl's mixing-length, One-Equation Models such as Spalart–Allmaras and Two-Equation Models such as, $k - \varepsilon$ and $k - \omega SST$.

The standard $k - \varepsilon$ model solves for the turbulent kinetic energy (k) and its rate of dissipation (ε). The modeled transport equation for k is shown in Equation (20), and for ε in Equation (21).

$$\frac{\partial(\rho k)}{\partial t} + \frac{\partial(\rho \bar{u}_j k)}{\partial x_j} = \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_k} \right) \frac{\partial k}{\partial x_j} \right] + P_k - \rho \varepsilon. \quad (20)$$

$$\frac{\partial(\rho \varepsilon)}{\partial t} + \frac{\partial(\rho \bar{u}_j \varepsilon)}{\partial x_j} = C_{\varepsilon 1} P_k \frac{\varepsilon}{k} - \rho C_{\varepsilon 2} \frac{\varepsilon^2}{k} + \frac{\partial}{\partial x_j} \left(\frac{\mu_t}{\sigma_\varepsilon} \frac{\partial \varepsilon}{\partial x_j} \right). \quad (21)$$

In these equations, k is turbulent kinetic energy, P_k signifies the production turbulent kinetic energy, ε denotes dissipation rate, σ_k indicates the turbulent Prandtl number for k , σ_ε , $C_{\varepsilon 1}$, $C_{\varepsilon 2}$ are model constants. The eddy viscosity μ_t is then calculated using k and ε , as shown in Equation (22),

$$\mu_t = \rho C_\mu \sqrt{k} L = \rho C_\mu \frac{k^2}{\varepsilon}. \quad (22)$$

Similarly, the $k - \omega$ model solves transport equations for k and the specific dissipation rate ($\omega = \varepsilon/k$). This model is often preferred for its better performance in near-wall regions and for boundary layer flows. The k contains a different dissipation term, i.e.,

$$\frac{\partial(\rho k)}{\partial t} + \frac{\partial(\rho \bar{u}_j k)}{\partial x_j} = P_k - \rho \beta^* k \omega + \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_k^*} \right) \frac{\partial \omega}{\partial x_j} \right], \quad (23)$$

and the modelled ω -Equation is

$$\frac{\partial(\rho \omega)}{\partial t} + \frac{\partial(\rho \bar{u}_j \omega)}{\partial x_j} = \alpha \frac{\omega}{k} P_k - \rho \beta \omega^2 + \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_\omega^*} \right) \frac{\partial \omega}{\partial x_j} \right]. \quad (24)$$

The eddy viscosity is expressed as:

$$\mu_t = \rho \frac{k}{\omega}. \quad (25)$$

Menter's Shear Stress Transport (SST) model is a popular variant blending $k - \omega$ near walls with the $k - \varepsilon$ model in the free stream to leverage the advantages of both.

2.5 Numerical Discretisation Approaches

Following the selection of a mathematical model, the subsequent crucial step involves choosing an appropriate discretisation method. This process entails approximating the governing differential equations through a system of algebraic equations defined for variables at a discrete set of spatial and temporal locations. Several methodologies exist for this purpose: the Finite Difference Method (FDM), the Finite Volume Method (FVM), and the Finite Element Method (FEM) representing the most prominent approaches. It is important to distinguish between domain discretisation (e.g., structured or unstructured meshing) and numerical discretisation (e.g., FDM, FVM, FEM), as they represent two independent but complementary layers in the computational setup. The following in-

depth descriptions of numerical discretisation approaches are derived from Ferziger et al. (2020).

2.5.1 Finite Difference Method (FDM)

The FDM proceeds by overlaying the solution domain with a grid structure. At each discrete grid point (node), the partial derivatives appearing in the PDE are replaced by finite difference approximations. These approximations are typically derived using Taylor series expansions or polynomial fitting, expressing the derivatives in terms of the function values at the node itself and a set of neighbouring nodes. This process transforms the original PDE into a system of algebraic equations, with one equation per grid node, where the nodal values of the dependent variables constitute the unknowns.

While FDM can, in principle, be formulated on various grid types, its practical application has been predominantly confined to structured grids. On such grids, the grid lines naturally align with local coordinate directions, simplifying the approximation of derivatives. FDM demonstrates effectiveness and simplicity on structured, regular grids, where developing higher-order accurate schemes is relatively straightforward. A typical structured grid exemplifying FDM is shown in Figure 26, where uniformly spaced nodes are arranged along Cartesian directions to facilitate finite difference approximations.




Figure 26. Cartesian grids commonly used in the FDM. Redrawn from (Ferziger et al., 2020).

2.5.2 Finite Volume Method (FVM)

The Finite Volume Method adopts a different starting point by using the integral form of the conservation equations, which directly expresses the conservation principle for a finite region of space.

In FVM, the computational domain is subdivided into a finite number of non-overlapping and contiguous control volumes (CVs). A computational node, where the dependent variable values are stored and computed, is typically located at the centroid of each CV. The integral form of the conservation law is then applied to each individual CV. This involves approximating surface integrals that represent fluxes of conserved quantities across control volume boundaries, and volume integrals that represent source or sink terms within the CV. Surface integrals are evaluated by expressing the variable values and their

gradients at the CV faces using interpolation based on the values at adjacent computational nodes. Suitable quadrature rules are used to approximate both surface and volume integrals. This procedure yields one algebraic equation per CV, implicitly linking the nodal value within that volume to the values in neighbouring CVs.

A significant strength of the FVM is its geometric flexibility; it can be readily applied to arbitrary grid types, including unstructured meshes, making it well-suited for complex geometries often encountered in practical simulations. The grid primarily serves to define the boundaries of the CVs and does not need to align with any specific coordinate system. Furthermore, the FVM is inherently conservative by construction, provided that the flux calculation across a shared boundary between two adjacent CVs is consistent, i.e., the flux leaving one CV is identical to the flux entering the neighbouring CV. Its conceptual basis, rooted directly in physical conservation principles applied to finite volumes, often makes it intuitive to understand and implement, contributing to its widespread use in engineering fields.

A primary challenge associated with FVM, particularly in three dimensions, is the development of schemes with accuracy higher than second-order. This difficulty arises from the multiple levels of approximation involved: interpolation of values to faces, potential differentiation to obtain gradients for fluxes, and numerical integration over surfaces and volumes. As illustrated in Figure 27, FVM commonly positions computational nodes at the centres of control volumes, allowing surface fluxes to be approximated at CV faces using interpolated values from adjacent nodes.




Figure 27. A cell-centred FVM grid. Control volumes are shown as squares with computational nodes (o) at their centres. The solid dots (•) define the boundary faces where boundary conditions are applied. Redrawn from Ferziger et al. (2020).

2.5.3 Finite Element Method (FEM)

The Finite Element Method shares similarities with FVM, particularly in its decomposition of the domain into discrete entities, termed finite elements. These elements typically form an unstructured mesh, commonly consisting of triangles or quadrilaterals in two dimensions, and tetrahedra or hexahedra in three dimensions.

The defining characteristic of FEM is its use of a weighted residual formulation. The governing PDEs are multiplied by a set of weight functions (also known as test functions), and the resulting equations are integrated over the entire computational domain (or, equivalently, element by element). Within each element, the solution field is approximated using basis functions (often called shape functions), typically polynomials, which are defined in terms of the solution values at the element nodes (usually vertices or other specific points on the element). These shape functions are constructed to ensure desired continuity of the solution across element boundaries. In the common Galerkin

FEM, the weight functions are chosen from the same family as the basis functions used for approximating the solution.

The approximation is substituted into the weighted integral form of the conservation law. The discrete algebraic equations are derived by requiring the weighted residual (the integral) to be zero for each weight function. This condition effectively minimises the error of the approximation in a weighted sense, selecting the best possible solution within the function space spanned by the chosen basis functions.

FEM offers significant advantages, most notably its high degree of geometric flexibility, making it adept at handling arbitrarily complex domains. An extensive body of knowledge exists regarding mesh generation techniques for FEM. Furthermore, element refinement, a key aspect of adaptive solution strategies, is straightforwardly implemented by subdividing elements. FEM possesses a strong mathematical foundation, allowing for rigorous error analysis and proofs of optimality for certain classes of PDEs.

A principal drawback, shared by methods employing unstructured grids (including FVM), relates to the structure of the resulting algebraic system. The matrices derived from FEM (and unstructured FVM) typically lack the regular banded structure found in FDM on structured grids. This less structured nature can pose challenges for developing highly efficient linear equation solvers, potentially increasing computational cost compared to structured-grid methods for problems where the latter are applicable.

3 Methodology

3.1 Experimental Setup and the Source of Data

Steady-flow benches (also known as steady-flow rigs or flow benches) have long been fundamental tools in internal combustion engine development. By applying a constant pressure differential across the intake or exhaust port and documenting the resultant flow rate and other flow attributes, they isolate the intrinsic flow capacity of intake and exhaust components without the confounding influence of a moving piston or fluctuating pressures (Yang et al., 2017). Kim et al. (2003) characterise the steady-flow bench as a “practical, powerful and widely used” tool that most engine manufacturers employ to form the initial design concept for a new engine. Steady-flow tests therefore underpin many comparative studies of port, valve and seat geometries, directly informing volumetric-efficiency and mixture-formation targets.

Three full-scale Wärtsilä 31 cylinder heads were tested at GKN Aerospace, Trollhättan in Sweden. The campaign served three overarching purposes that are central to this thesis, namely, product improvement, method development and generation of validation data. To meet these goals, three complementary test modes were executed, as described in Table 4.

Table 4. Flow bench test modes for full-scale W31 cylinder heads: static (0 – 100 % lift sweep), scavenging (\approx 5 – 20 % lift with both valves open), and dynamic (realistic moving lift profiles).

Test mode	Valve configuration	Typical lift range	Purpose
Static	The valves corresponding to the given configuration (e.g., intake or exhaust) are swept from barely open to full lift while the counterpart valves remain closed.	0–100 % of design lift	Base flow coefficients; isolate individual valve behaviour.

Test mode	Valve configuration	Typical lift range	Purpose
Scavenging	Intake and exhaust valves are held simultaneously at low lifts representative of the scavenging phase; a dummy piston blocks the liner; accordingly, the air exits immediately through the exhaust.	≈ 5–20 %	Assess cross-flow interaction during early/late gas exchange.
Dynamic	Each valve is driven through realistic opening and closing profiles by a hydraulic valve train actuator; different durations and pressure ratios are imposed.	Profile-dependent	Capture transient effects.

These objectives and test types provide the experimental backbone for the numerical work done in this thesis, and CFD validation draws exclusively on the static flow bench data.

3.1.1 Physical Setup

Compressed air for the flow bench campaign is supplied by an industrial compressor and stored in an underground reservoir that is partially water-filled. The incompressible water column minimises gas-volume fluctuations, allowing for the high-pressure ratios required to replicate engine conditions. Before reaching the test rig, the air passes through a dryer to ensure a low dew-point and eliminate moisture-induced density errors.

On entering GKN Aerospace's test bay the flow encounters a calibrated sharp-edged orifice plate, across which a differential-pressure transmitter provides the primary mass-flow measurement. The stream then enters Wärtsilä's inlet pipework, where a Pitot-static probe supplies an independent cross-check of the mass flux and total pressure.

Downstream of the inlet pipe, the air can follow two alternative routes, selected by the operator according to the test mode, either exhaust flow configurations or intake.

In the exhaust flow configuration, the stream enters the bottom of the cylinder liner, passes the exhaust valves and exits through the exhaust port into a wall-mounted flange before discharging to atmosphere as it can be observed in Figure 28.




Figure 28. Exhaust flow setup. In this configuration, the air enters Wärtsilä's inlet pipe flowing into the bottom of the cylinder liner exiting passed the exhaust valves out through the exhaust port then through the Waste Gate Valve and finally discharged through the outlet. Proprietary components are intentionally blurred in accordance with Wärtsilä confidentiality requirements.

In the intake flow configuration, the stream rises through a riser section, passes the inlet port and the intake valves, then leaves the cylinder via the liner bottom and the same wall flange as shown in Figure 29.




Figure 29. Air enters the receiver section downstream of the inlet pipe, rises through the riser pipes, crosses the intake port and valves, then descends the cylinder liner and exits through the Waste Gate Valve to the outlet stack. Proprietary components are blurred in compliance with Wärtsilä confidentiality.

The scavenging arrangement, although executed during the experimental campaign, is outside the scope of this thesis and is therefore not treated further. Figure 30 demonstrates the full-scale flow bench setup at GKN Aerospace.




Figure 30. Schematic diagram of a pressurised flow bench system for cylinder head testing, showing airflow from the air source through the orifice plate entering inlet pipe equipped with a Pitot-static probe. It then rises through the cylinder head and finally discharges to atmosphere via the exhaust port. Arrows indicate flow direction. Drawn by the author.

This schematic arrangement is realised in the physical test rig at GKN Aerospace, as photographed in Image 1. The image demonstrating the exhaust flow configuration, allows for direct correlation with the schematic's components. The large-diameter pipework labelled 'Inlet' on the right of the image corresponds to the inlet section downstream of the orifice plate in Figure 30; the component identified as the 'Wärtsilä pitot static Tube' in the schematic is marked as the 'Wärtsilä's MAF sensor' in the physical setup. This pipe-work feeds the 'Flow bench', installed on the blue-framed structure known as 'bench surface' that supports the cylinder head assembly for flow characterisation. After passing through the cylinder head, the air exits via the exhaust port and is directed towards the 'Outlet' on the left. The photograph also reveals additional hardware not present in

the simplified schematic, such as the Waste Gate (WG) valve installed for pressure regulation in the exhaust line.



Image 1. Exhaust flow configuration at the GKN Aerospace facility. Air enters via the 'Inlet' (right), passes through the cylinder head mounted on the central test stand (labelled 'Flow bench'), and exits through the 'WG valve' to the 'Outlet' (left). Proprietary components are blurred to maintain confidentiality.

3.1.2 Design of Experiments

The experimental data was systematically gathered by Wärtsilä experts through a structured test campaign involving three full-scale Wärtsilä 31 cylinder heads (CH1, CH2, and CH3). These cylinder heads differ in valve seat geometry, the number of installed valves and their respective sizes representing variations of interest for flow behaviour analysis.

CH1 served as the experimental baseline, representing the standard configuration currently in field service. It features distinct seat angles for the intake and exhaust valves, along with chamfered valve seat pockets.

CH2 was a design variation focused on valve seat geometry. In contrast to the baseline, CH2 was configured with symmetric valve seat angles, utilising an identical proprietary profile for both intake and exhaust. Furthermore, the valve seat pockets were rounded.

CH3 was a specialised configuration originally developed for research on a hydrogen-fuelled single-cylinder engine. Its defining feature is a five-valve architecture, comprising three intake and two exhaust valves. It features the same asymmetric valve seat geometry as the CH1 baseline.

In order to study the flow behaviour in these three CHs, a detailed test matrix was executed for both intake and exhaust configurations, as summarised in Tables 5 and 6. The primary test was the 'Valve sweep,' performed at several key pressure ratios i.e., 1.1, 1.5, 2.0 for exhaust and 1.1, 1.25 and 1.5 for intake. During each sweep, the valve lift was varied from its minimum to maximum value in 15 to 17 discrete steps, resulting in 15 to 17 data points per run. Each target pressure ratio was also tested using multiple combinations of absolute inlet and outlet pressures to assess any influence of the absolute pressure level. To further analyse the effects of pressure, pressure ratio sweep tests were also performed at fixed mid and high lift positions. Additional tests, such as single-valve-open and flow through single riser pipe, were included to investigate flow symmetry and individual valve contributions. In total, the campaign generated ~1500 test points.

Across all tests, several key quantities were measured for each data point. Mass flow rate was the primary measurement, determined via an orifice plate and cross-checked with a Pitot-static tube. Extensive pressure and temperature readings were also taken at multiple locations, including the inlet, outlet, and on the cylinder head itself.

Table 5. Intake flow bench test matrix.

Test Type	Cylinder Heads	Description	Pressure Ratio Range	Normalised Lift L/L_max. Range (0-1)
Valve sweep	CH1, CH2, CH3	Valve lift sweep with three distinct inlet/outlet pressures	1.1, 1.25, 1.5	0.043 to 1.000

Test Type	Cylinder Heads	Description	Pressure Ratio Range	Normalised Lift L/L_max. Range (0-1)
Pressure ratio sweep at mid-lift	CH1, CH2, CH3	Fixed lift at 0.314, pressure ratio sweep	1.1 to 8.9	0.314
Static A valve open	CH1, CH2, CH3	Only one valve open during lift sweep	1.5	0.043 to 1.000
Static B valve open	CH1, CH2, CH3	Alternate single-valve configuration	1.5	0.043 to 1.000
180° rotated valve configuration	CH1, CH2, CH3	Symmetry-rotated valve orientation during lift sweep	1.5	0.043 to 1.000
Flow through one riser pipe	CH1, CH2, CH3	One riser pipe active, others blocked	1.5	0.043 to 1.000

Table 6. Exhaust flow bench test matrix.

Test Type	Cylinder Heads	Description	Pressure Ratio Range	Normalised Lift Range
Valve sweep	CH1, CH2, CH3	Valve lift sweep with three distinct inlet/outlet pressures	1.1, 1.5, 2.0	0.029 to 1.0
Pressure ratio sweep at high and mid-lift	CH1, CH2, CH3	At lift 1.0: pressure ratio of 1.1–8.0; at lift 0.314: pressure ratio of 1.1–2.9	1.0 to 8.0	0.314 or 1.0
Static C valve open (CH1, CH2)	CH1, CH2	One valve open during lift sweep	1.5	0.029 to 1.0
Static C valve open (CH3)	CH3	Valve D open	1.5	0.029 to 1.000
Static D valve open (CH1, CH2)	CH1, CH2	Alternate valve open during lift sweep	1.5	0.029 to 1.000
Static D valve open (CH3)	CH3	Valve E open	1.5	0.029 to 1.000
180° rotated valve configuration	CH1, CH2, CH3	Symmetry-rotated valves during lift sweep	1.5	0.029 to 1.000

Test Type	Cylinder Heads	Description	Pressure Ratio Range	Normalised Lift Range
Flow through one riser pipe	CH1, CH2, CH3	Single riser pipe flow path	1.5	0.029 to 1.000

3.1.3 Uncertainty Analysis for Experimental Results

Experimental data without quantified uncertainties cannot be meaningfully compared, validated, or interpreted. According to Moffat (1988), an uncertainty band defines the range within which the true value is believed to lie at stated odds and therefore allowing both author and reader to judge significance consistently. Therefore, a quantitative statement of the associated uncertainties is imperative to assess the confidence level of the results derived from the flow bench campaign. While a formal, sensor-by-sensor uncertainty analysis was not performed, a representative single-sample uncertainty analysis was conducted to establish a credible, baseline estimate of the measurement system's reliability. This analysis focuses specifically on the Wärtsilä ITABAR Pitot-static tube, as the manufacturer's detailed sizing sheet provided the necessary component specifications and nominal values required for a full propagation of uncertainty calculation. The analysis is based on the manufacturer's design specifications for the exact equipment, centered on a nominal operating point.

The methodology employed is the single-sample uncertainty analysis, a framework developed specifically for experiments where each test point is typically run only once. This approach is ideally suited for the broad parametric sweeps conducted in this flow bench campaign. The approach tracks the propagation of individual measurement errors to the final calculated result using the root-sum-square (RSS) method. In this method, a final result, R , is calculated from a set of independent measured variables (X_1, X_2, \dots, X_n), the overall uncertainty in the result, δR , is determined by:

$$\delta R = \left\{ \sum_{i=1}^N \left(\frac{\partial R}{\partial X_i} \delta X_i \right)^2 \right\}^{1/2}. \quad (26)$$

In Equation 26, δR is the total uncertainty, δX_i is the uncertainty interval associated with each individual measurement, $\frac{\partial R}{\partial X_i}$ is the sensitivity coefficient, representing the partial derivative of the result R with respect to the input variable X_i . In accordance with standard engineering practice (Moffat, 1988), a 95% confidence interval is used for this analysis.

Estimation of individual uncertainties (δX_i) for each primary measurement is itself a combination of two distinct error types: bias errors and precision errors. Bias errors are fixed, systematic offsets that remain constant during a test run. Sources include instrument calibration inaccuracies, and inherent limitations in the data acquisition hardware. And precision errors are non-systematic, random fluctuations that cause scatter in repeated measurements. They arise from sources such as electronic noise in the instrumentation, small thermal drifts, and process unsteadiness. The total uncertainty for a single measurement, δX_i , is the RSS combination of its bias (B_i) and precision (P_i) components:

$$\delta X_i = (B_i^2 + P_i^2)^{1/2}. \quad (27)$$

Since direct calibration data was unavailable for the specific installation, the uncertainty analysis uses values derived from the ITABAR sizing sheet (Order No. 24/52013) as representative estimates. These values correspond to a nominal operating condition similar to the flow bench setup but are not exact duplicates of the test configuration. Here are the system specifications and nominal data:

- Flow Sensor: Intra-Automation ITABAR IBR-26
- Differential Pressure (DP) Transmitter: Aplisens APR-2000ALW (calibrated range: 0 to 302.34 mbar)
- Nominal Data Point:
 - Mass Flow Rate (\dot{m}_0): 8.00 kg/s
 - Differential Pressure (ΔP): 193.17 mbar (19,317 Pa)
 - Line Pressure (P_1): 8.50 bar(g) (951,330 Pa absolute), assuming atmospheric pressure

- Line Temperature (T_1): 18.00 °C (291.15 K)

Based on the propagation of uncertainties from the primary instruments and mechanical tolerances (detailed in Appendix 2), the calculated total uncertainty in the mass flow rate at the nominal design point is: $\dot{m} = 8.00 \pm 0.131$ kg/s. This corresponds to a relative uncertainty of ± 1.6 %.

A sensitivity analysis revealed that the total uncertainty is overwhelmingly dominated by two factors: the manufacturing tolerance of the pipe diameter (72.89% contribution) and the calibration uncertainty of the ITABAR K-Factor (24.52% contribution). The combined uncertainty from the electronic pressure, differential pressure, and temperature transmitters accounts for slightly more than 2% of the total.

The single-sample analysis provides a baseline uncertainty of ± 1.6 %, derived from manufacturer specifications. However, according to Figiola & Beasley (2020), standard engineering practice requires acknowledging elemental error sources beyond instrument ratings that can introduce considerable bias. In this experiment, several unquantified systematic uncertainties exist.

The detailed uncertainty analysis presented in this thesis is focused on the ITABAR Pitot-static tube for a specific reason: it was the only primary mass flow sensor for which a complete manufacturer's specification sheet (the Winflow sizing sheet) was available. This detailed data enabled a full, bottom-up propagation of uncertainty calculation using the RSS method. Even with this detailed sheet, it is important to note that the uncertainties for some input variables, such as absolute pressure and temperature, were not explicitly provided for this installation. In these cases, reasonable engineering estimates were made based on the specifications of standard, high-quality industrial instruments, as detailed in Appendix 2.

An equivalent, detailed analysis was not performed for the GKN orifice plate due to the lack of a comparable calibration envelope or detailed specification sheet. Regarding the other primary measurements, such as inlet pressure, outlet pressure, and temperature, their specified instrument uncertainties are directly incorporated as inputs into the main mass flow uncertainty calculation. They are not treated as standalone validation benchmarks themselves, but rather as critical boundary conditions for the CFD simulations. This approach concentrates the most rigorous analysis on the primary performance metric used for validation, i.e., the mass flow rate. In this study, the measurements from both the Wärtsilä and GKN sensors serve as the key experimental benchmarks for the CFD results.

Moreover, the ITABAR-Sensor Type IBR Operating Manual (n.d.) explicitly states that the validity of its K-factor is contingent on an 'undisturbed' flow profile, typically requiring an extended straight-pipe run. Deviations from this ideal condition in the physical layout can lead to systematic errors in the mass flow rate. Other installation errors, such as sensor misalignment ($<5^\circ$ tolerance) or improper orientation, contribute additional bias.

Consequently, the $\pm 1.6\%$ figure should be interpreted as the best-case uncertainty, representing the fundamental limit imposed by the instrumentation and mechanical tolerances under ideal conditions. The true uncertainty, encompassing these potential installation effects, is likely to be larger. The impact of these factors will be explored in the discussion of the results in Chapter 4.

3.2 Numerical Setup

3.2.1 Computational Mesh Generation

The process begins with preparing the Computer-Aided Design (CAD) geometry, from which the fluid domain is defined. This domain is subsequently discretised into a

computational mesh for the numerical analysis. Meshes for CH1 in both the configurations have been generated with SnappyHexMesh (SHM). SHM is an automatic mesh generation tool in the open-source OpenFOAM CFD suite. It produces predominantly hexahedral and split-hexahedral (polyhedral) meshes by “sculpting” an initial block-structured grid to fit complex geometries. Starting from a simple background mesh (e.g. from OpenFOAM’s blockMesh), SHM iteratively refines cells near surface features and then snaps the mesh to align with the geometry (Soner & Ozturan, 2015). This approach yields hexahedral-dominant unstructured grids, offering a compromise between fully structured hexahedral meshes and tetrahedral meshes.

Figure 31 presents the computational mesh for the CH1 intake case, generated using SHM. The analysed fluid domain consists of the riser pipes, intake port, cylinder head, and cylinder liner, with the cylinder bottom defining the outlet patch.




Figure 31. Computational mesh of the intake flow CH1 fluid domain generated with SHM (left), and a cross-sectional view of the intake valves region (right).

One of SHM’s goals is to ensure acceptable mesh quality metrics (e.g. cell non-orthogonality, skewness) despite the automatic nature of the algorithm (*A Comprehensive Tour of SnappyHexMesh*, 2012). By refining a Cartesian mesh and cutting cells against the surface, SHM generates a hexahedral-dominant mesh that tends to require fewer total cells than an equivalent tetrahedral mesh for the same geometry. Montorfano (2017) noted that a cut-cell mesh from SHM can use up to 40% fewer elements than a purely

tetrahedral mesh for the same geometry, owing to the efficiency of hex/polyhedral shapes. This cell count reduction, combined with the generally higher numerical fidelity of aligned hex cells, can improve solution accuracy compared to tetrahedral grids. The mesh quality metrics for the flow bench simulations in this work, summarised for two representative CH1 configurations in Table 7, demonstrate these advantages.

Table 7. Mesh quality metrics for the CH1 intake and exhaust configurations at representative low (0.14) and full (1.0) valve lifts, normalised by the maximum lift (L/L_{max}).

Case	Cells [$\times 10^6$]	Hexahedra [%]	Polyhedra [%]	Max non-orth ($^\circ$)	Max skew	High skew faces	Status
Intake 0.14	3.19	79.9	17.5	64.9	3.93	0	OK
Intake 1.0	3.19	80.0	17.4	65.0	4.70	1	1 skew warning
Exhaust 0.14	2.24	80.4	16.6	64.9	9.62	6	1 skew warning
Exhaust 1.0	2.24	80.6	16.4	64.9	6.33	4	1 skew warning

Table 8 presents the mesh quality report for the CH1 configurations, as generated by the `checkMesh` utility. The results show that the meshes meet the key quality thresholds for hex-dominance and non-orthogonality, with only a small number of high-skewness cells detected in the exhaust cases as explained by Greenshields (2024).

Table 8. Mesh-quality scorecard with key take-aways.

Metric & Threshold	Intake 0.14	Intake 1.0	Exhaust 0.14	Exhaust 1.0	Verdict
Hexahedral dominance $\geq 75\%$	80 %	80 %	80 %	81 %	All cases meet the efficiency target.
Max non-orthogonality $\leq 70^\circ$	64.9 $^\circ$	65.0 $^\circ$	64.9 $^\circ$	64.9 $^\circ$	Safely below the RANS-solver comfort zone.

Metric & Threshold	Intake 0.14	Intake 1.0	Exhaust 0.14	Exhaust 1.0	Verdict
Max skewness ≤ 5 for robust RANS	3.93	4.70	9.62	6.33	Acceptable on intake; exhaust exceeds guideline but affects < 0.0003 % of cells.
Failed mesh checks	0	1(skew)	1(skew)	1(skew)	Sparse high-skew faces flagged; treat with conservative numerics.

However, challenges remain in complex cases. SnappyHexMesh can struggle to retain high quality in very intricate geometries or when extremely fine near-wall resolution is required. For example, an automotive aerodynamic study at Audi found that SHM “is unable to generate low- y^+ grids of sufficient quality for [production car] geometries.” In other words, when targeting a very fine boundary layer mesh ($y^+ \approx 1$) on a complicated vehicle surface, the algorithm could not produce a grid without quality defects (e.g. collapsed or highly skewed cells) (Ashton et al., 2018). Similar observations have been made in this thesis during intake flow simulations where meshes collapsed at normalised valve lifts below 0.14.

However, despite its limitations, it offers further advantages. One advantage is minimal manual effort; defining a SHM dictionary (meshing parameters) is the primary task, and once set, the mesh is generated automatically without user intervention. Montorfano (2017) noted that the “time to prepare the case is negligible” compared to the actual meshing runtime. Another major strength of SHM is its full parallelisation, which allows large meshes to be generated on HPC (High-Performance Computing) clusters. All three stages of SHM (castellated refinement, snapping, layering) are implemented to run on a decomposed domain, communicating between processors as needed. In fact, SHM by default runs in parallel if the OpenFOAM case is decomposed, and it can use any number of processors with no additional effort by the user. Therefore, considering the automation, time-saving that it offers and parallelisation capabilities, it has been opted for this study.

3.2.2 Case Setup and Input/Output Structure

All numerical flow bench simulations were performed using the open-source CFD code, OpenFOAM-12. A standard OpenFOAM case structure is employed, which logically separates initial and boundary conditions (`0/` directory), time-invariant physical properties and mesh data (`constant/` directory), and solver control settings (`system/` directory). In a transient compressible PIMPLE case (solver run via `foamRun -solver fluid`), these files work together to define the problem.

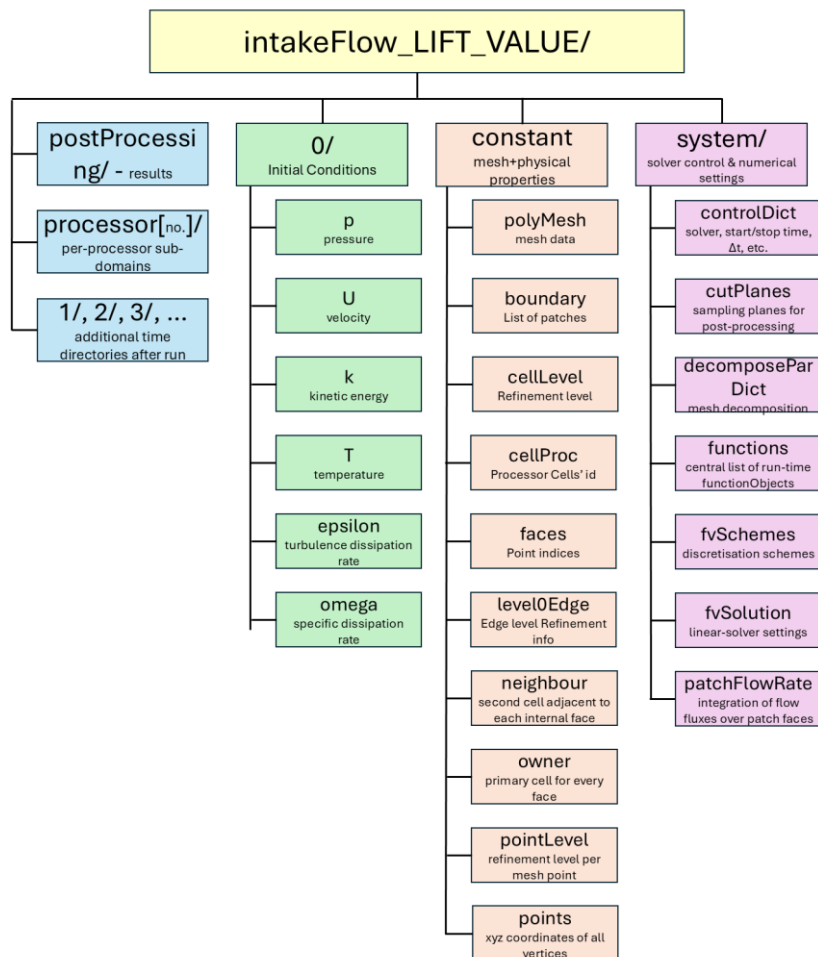


Figure 32. Directory structure of a typical OpenFOAM simulation case used in this thesis. Core folders define initial conditions (`0/`), physical properties (`constant/`), and solver settings (`system/`). Output data are stored in `postProcessing/`, time directories, and parallel subdomains (`processor[n]/`).

0 (Initial/Boundary Conditions)

The (0/ directory) contains the initial conditions for the entire computational domain at time $t = 0$. For this transient, compressible RANS simulation, the primary transported variables are pressure (p), velocity (U), temperature (T), and the turbulence quantities required by the chosen model. For example, for the $k - \omega$ SST model turbulence variables are turbulent kinetic energy (k) and specific dissipation rate (ω).

These files define the initial state of the fluid in every one of the mesh's volumetric cells. Therefore, the data is stored at the highest possible spatial resolution, a unique value for each field (p, U, T , etc.) is assigned to every cell centre in the domain. The physical relevance of these fields is tied directly to the governing equations of fluid dynamics (as outlined in Section 2.3). p, U , and T are essential for solving the conservation of mass, momentum, and energy, while k and ω are required to close the RANS equations and model the effects of turbulence.

The boundary conditions, specified for each patch within these files, are critical for defining how the domain interacts with its surroundings. The choices outlined in Table 9 were made in the attempt to accurately represent the physical flow bench apparatus. For instance, the use of `omegaWallFunction` is crucial for correctly modelling the turbulent boundary layer near solid walls without requiring an impractically fine mesh resolution right down to the viscous sublayer. The below Table 9 serves as a detailed reference describing every file present in 0 directory.

Table 9. Numerical Flow bench boundary condition reference table

Patch / Regex	Boundary Condition Type	Physical Meaning	Why Chosen
U			

Patch / Regex	Boundary Condition Type	Physical Meaning	Why Chosen
(iv ev liner intake-port exhaustport flame-plate piston)	noSlip	$u = 0$ at solid wall	Viscous metal walls; captures boundary-layer losses in the continuum regime
inlet	pressureInletOutletVelocity	Inflow velocity from Δp , zero-grad on outflow	Handles reverse flows; robust for transient port
outlet	pressureInletOutletVelocity	Same as inlet	Ensures consistent two-way treatment
ncc_*	slip	Zero normal velocity, free tangential	Non-computational or symmetry interior planes
P			
(iv ev liner intake-port exhaustport flame-plate piston)	zeroGradient	$\partial p / \partial n = 0$ at wall	Impermeable; avoids spurious pressure peaks
inlet	uniformTotalPressure (ramped)	Prescribed total pressure $p_0(t)$	Models gradual pressurisation
outlet	uniformTotalPressure (constant)	Fixed downstream p_0	Sets reference pressure
ncc_*	zeroGradient	Symmetry/outflow	Non-computational. Does not disturb pressure field
T			

Patch / Regex	Boundary Condition Type	Physical Meaning	Why Chosen
(iv ev liner intake-port exhaustport flame-plate piston)	zeroGradient	Adiabatic wall ($\partial T / \partial n = 0$)	No combustion; Also, unknown wall-temperature values from experiments
inlet	fixedValue	Sets inflow temperature	Reservoir-conditioned air
outlet	zeroGradient	Convective outflow	Lets solver determine exiting temperature
ncc_*	zeroGradient	Symmetry	—
ω			
(iv ev liner intake-port exhaustport flame-plate piston)	omegaWallFunction	Log-law formula for ω near wall	Correct near-wall scaling for $k - \omega/SST$
inlet	turbulentMixingLengthFrequencyInlet	ω from mixing length ℓ	Specify intuitive turbulence scale
outlet	turbulentMixingLengthFrequencyInlet	Same as inlet (mirrored)	Stable during back-flow
ncc_*	zeroGradient	Symmetry	—
ϵ			
(iv ev liner intake-port exhaustport flame-plate piston)	epsilonWallFunction	Wall log-law relation for ϵ	Standard $k - \epsilon$ wall treatment

Patch / Regex	Boundary Condition Type	Physical Meaning	Why Chosen
inlet	turbulentMixing- LengthDissipation- RateInlet	ε from mixing length ℓ	Consistent with specified turbulence
outlet	turbulentMixing- LengthDissipation- RateInlet	Same as inlet	Handles reversal
ncc_*	zeroGradient	Symmetry	—
<i>k</i>			
(iv ev liner intake- port ex- haustport flame- plate piston)	kqRWallFunction	Models the near wall region using the log law	Standard wall treatment for such cases.
inlet	turbulentIntensi- tyKineticEnergyIn- let	Specifies k based on the turbulence intensity (I).	Provides a realistic turbulence level based on measured experimental data.
outlet	turbulentIntensi- tyKineticEnergyIn- let	Same as inlet	Supports back-flow
ncc_*	zeroGradient	Symmetry	—

Constant directory

It defines the physical context of the problem. Unlike the field data in the (0/ directory) which is resolved on a per-cell basis, the files here specify global physical laws, material properties, and the computational grid itself; these inputs stay the same during the run. Because these files are time-invariant, OpenFOAM reads them once at start-up and keeps them fixed throughout the transient simulation.

The contents of this directory are fundamental as they establish both the geometric space and the physical rules governing the fluid's behaviour. The “polyMesh” subdirectory is arguably the most critical component, as it contains the discretised geometry and defines the ultimate spatial resolution of the entire simulation. The quality and topology of this mesh directly influence the accuracy of all calculated gradients and fluxes.

The other key files dictate the modelling framework. The “momentumTransport” file selects the turbulence model and its empirical constants, a choice critical for capturing the effects of turbulence in a computationally feasible manner for high-Reynolds number port flows. The “physicalProperties” defines the thermodynamic behaviour of the fluid (air), which is essential for a compressible solver where density and viscosity change significantly with pressure and temperature.

Finally, the custom “FlowbenchDict” serves as a centralised control dictionary. It allows key simulation parameters, such as valve lift and boundary conditions, to be defined in a single location, which enables the systematic and automated variation of cases. This functionality is central to the PyFlowBench workflow. The choices made in this directory, detailed in Table 10, represent the core modelling assumptions of this study.

Table 10. Main dictionaries in constant/

File / Directory	What it holds	Why it matters for this case
constant/polyMesh (directory)	Mesh defined by points, faces, connectivity & patch list.	Defines the intake or exhaust geometry and cell topology the solver needs for fluxes, gradients and boundary-condition mapping
constant/momentumTransport	Turbulence framework: simulation-Type RAS, Turbulence Model, and the empirical coefficients for $k - \omega$ SST and $k - \epsilon$ RANS models	Selects and configures the two-equation RANS closure model appropriate for high-Re, compressible port flow and allows for turbulence model selection.
constant/physicalProperties	Thermodynamic & transport model for air: perfect-gas Equation of State, Constant Specific	Provides $\rho(p, T)$ and $\mu(T)$ to the compressible PIMPLE solver so density and viscosity vary realistically with pressure/temperature

File / Directory	What it holds	Why it matters for this case
	Heat $(Cp) = 1005 J kg^{-1} K^{-1}$, Sutherland viscosity law	
constant/Flow benchDict	Acts as the central controlling centre for the entire simulation. Allows for imposing the case specific lift, Boundary conditions, initial turbulence values, directory paths and macro definitions and etc	Keeps all bench-specific settings in one file; macros propagate into 0/ and system/ dictionaries for easy case re-configuration required by PyFlowBench package.

System Directory

If the (constant/) directory defines the physical problem, the (system/) directory provides the numerical instruction manual for how to solve it. Its dictionaries govern the execution of the solver, dictating how the governing equations are discretised, linearised, and marched forward in time. These settings do not change the underlying physics but are critical for ensuring the numerical solution is stable, accurate, and obtained in a reasonable amount of time.

This directory provides fine-grained control over the solution process. The “controlDict” file manages the temporal control, defining the simulation's start and end times, the time-step (Δt), and the frequency of data output. The use of an automatic time-step adjusted by the courant number is crucial for maintaining stability in this compressible flow case where velocities can vary dramatically.

The “fvSchemes” dictionary controls the mathematical accuracy of the simulation by specifying the numerical schemes used to approximate the spatial and temporal derivatives in the governing equations. The selections here represent a deliberate trade-off between accuracy (favouring second-order schemes) and stability (requiring limiters to prevent non-physical oscillations in high-gradient regions like the valve jet).

The “fvSolution” file controls the algorithmic convergence, specifying the linear solvers and tolerance for each discretised equation and managing the PIMPLE algorithm, which

is essential for robustly coupling the pressure and velocity fields in a transient, compressible flow. Together, these settings directly influence the simulation's fidelity and the magnitude of the numerical uncertainties discussed in Section 3.2.3. Table 11 summarises the core dictionaries that orchestrate the solver's behaviour to meet the specific challenges of these flow bench simulations.

Table 11. Core dictionaries in system/ directory used in flow bench simulations conducted in OpenFOAM.

File / Dictionary	Key contents	Role in this case
system/controlDict	Run control parameters: start/stop time, Δt , automatic time-step, write interval; loads custom cutting-plane sampling via includeEtc surface.cfg	Governs the PIMPLE solver execution, defining the simulation timeline, time-step control, and the frequency of data output and run-time sampling.
system/fvSchemes	Finite-volume schemes: Euler ddt, limitedLinear & upwind convection, limited Laplacian, cell-limited gradient	Chooses stable and accurate discretisation for the high-speed, compressible port flow
system/fvSolution	Linear-solver settings and PIMPLE algorithm controls.	Ensures pressure–velocity coupling converges each time step even with transonic pockets
system/decomposeParDict	numberOfSubdomains accordingly; decomposer/distributor scotch	Automates balanced domain decomposition based on the chosen number of cores
system/functions	#includeFunc cutPlanes; patch flow-rate monitor; adaptive Δt table; probes & field extremes	Runs in-solver diagnostics (mass-flow, slices to VTK, sensor probes) so no separate post-step is needed

Data Extraction and Post-Processing

While the (0/, constant/, and system/) directories define the simulation inputs, the primary goal is to extract meaningful outputs for analysis. These outputs are generated during the simulation run and stored in two main locations: (time-step) directories and the (postProcessing/) directory.

As the simulation progresses, OpenFOAM writes the complete solution fields (p , U , T , k , ω , etc.) to disk at specified intervals, creating new directories named after the simulation time (e.g., 1/, 2/, 3/). The frequency of this output is controlled by the “writeInterval” parameter in (system/controlDict). While this data provides a comprehensive, high-resolution snapshot of the entire flow field for detailed 3D visualisation (as presented in Section 4.4), storing it at every time-step is computationally prohibitive due to large file sizes. Therefore, it is saved at a lower temporal frequency, sufficient for qualitative analysis and understanding key flow phenomena like jetting, flow separation, and the formation of swirl and tumble structures.

For quantitative analysis and direct comparison with experimental data, targeted data is extracted at a much higher temporal frequency using OpenFOAM's run-time “functionObjects,” which are defined in (system/functions). The outputs from these functions are stored in the (postProcessing/) directory. For this thesis, the key quantities of interest are:

- **Mass Flow Rate:** This is the primary performance metric for a flow bench. It is calculated by integrating the flux of density and velocity over the inlet or outlet patches at every time-step using the “patchFlowRate” function. This provides a high-resolution time-series of the mass flow, which is essential for comparing simulation results directly against the experimental data discussed in Section 4.2 and 4.3.
- **Temperature probe:** For quantitative validation against experimental measurements, the OpenFOAM probes functionObject was employed. A probe was defined at the virtual coordinates corresponding to the physical temperature sensor location within the cylinder head. This function samples the temperature field (T) at this discrete point with high temporal frequency, i.e., every time-step, generating a time-series output file. To obtain a stable metric for comparison, the simulated temperature from this time-series is time-averaged once the flow reaches a quasi-steady state. This value is then directly compared with the experimental

reading to validate the thermal model's accuracy, which is critical as temperature directly influences fluid density and, consequently, the predicted mass flow rate.

- **Cut planes for Flow Structure Analysis:** To analyse coherent in-cylinder flow structures, which are precursors to combustion-critical swirl and tumble motion, the surfaces sampling function was utilised. This function extracts specified field data (e.g., U, p, T) onto predefined two-dimensional geometric planes at regular time intervals during the simulation. These planes are strategically placed, for instance, in a cross-tumble orientation below the intake valves. The resulting 2D data sets enable a qualitative visualisation of flow phenomena (e.g., valve jets, separation zones). This targeted data extraction provides detailed insight into the evolution of flow structures without the prohibitive computational cost of storing and processing the full 3D domain data at a high temporal resolution.

By using this dual approach, saving full-field data infrequently and key performance metrics frequently, we can conduct a thorough analysis that is both computationally efficient and physically insightful.

3.2.3 Numerical Uncertainty Analysis

All numerical simulations are, by their nature, approximations of physical reality. Therefore, a credible CFD study requires a systematic analysis and quantification of the errors inherent in the solution process. This process, often referred to as Verification and Validation (V&V), is essential for establishing confidence in the simulation results before they can be meaningfully compared with experimental data. The methodology for this assessment is based on the widely accepted framework which categorises numerical errors into three primary sources: modelling, discretisation, and iteration errors (Ferziger & Perić, 2020).

A crucial aspect of this framework is the hierarchical order in which these errors are assessed. As recommended by Ferziger & Perić (2020), the analysis should proceed in the

reverse order of how the errors are generated. First, iteration errors must be shown to be negligible to ensure the algebraic equations are solved correctly. Second, discretisation errors are quantified or controlled to ensure the solution is not unduly influenced by the grid. Only when these numerical errors are controlled can one proceed to the final step of validation, where the remaining difference between the simulation and experimental data can be more confidently attributed to the inherent assumptions of the physical models themselves. This section details the practical steps taken to address each of these error sources in turn.

Control of Iteration Error

Iteration error, also known as convergence error, is the residual error remaining from the numerical solution of the linearised algebraic equations within a single time step. It is defined as the difference between the exact solution of these discrete equations and the approximate solution obtained after a finite number of solver iterations (Ferziger et al., 2020). If not properly controlled, this error can accumulate and corrupt the solution, making it impossible to assess other, more physically significant error sources. In this study, a multi-faceted strategy was implemented within the “system/fvSolution” dictionary to ensure iteration errors were minimised to a level significantly below other sources of uncertainty. This strategy is principally managed through the PIMPLE algorithm and the settings of the linear solvers.

The PIMPLE algorithm is a hybrid solver framework available in OpenFOAM that merges the logic of two classic algorithms: PISO (Pressure-Implicit with Splitting of Operators) and SIMPLE (Semi-Implicit Method for Pressure Linked Equations). The PIMPLE algorithm allows for a blend of these two approaches. The specific settings used in this work are configured to operate primarily in an efficient, PISO-like mode, controlled by the following parameters:

- `nOuterCorrectors 1`: This parameter dictates how many times the entire system of equations is assembled and solved per time step. By setting it to 1, the

simulation operates in a PISO-like mode, which is computationally efficient and well-suited for the small time steps used in this transient analysis.

- `nCorrectors 2`: Within the single outer loop, this parameter specifies that two "PISO corrector" loops are performed. These loops iteratively correct the pressure and velocity fields to ensure they satisfy the conservation of mass and momentum. Using two correctors provides a good balance between computational cost and the accuracy of the pressure-velocity coupling.
- `nNonOrthogonalCorrectors 2`: Real-world geometries, such as the intake port, necessitate meshes with some degree of non-orthogonality. This can introduce numerical errors when calculating gradients across cell faces. This parameter adds two explicit correction steps to account for and remove these errors, thereby improving the solution's accuracy on the practical, complex mesh used in this study.

Also, within each of the correction loops described above, a linear solver is called to solve the large system of algebraic equations for each variable (p , U , k , etc.). For most variables, a Preconditioned Bi-conjugate Gradient Stabilised (PBiCGStab) solver was selected, an efficient iterative method well-suited for the matrices produced in CFD. The convergence of this solver is governed by a dual-tolerance strategy:

- An absolute tolerance is the fixed target residual value the solver must reach. In this case, it was set to $1e-9$.
- A relative tolerance (`relTol`) instructs the solver to stop if the residual has been reduced by a certain factor relative to its starting value. This was set to 0.001 (a reduction of three orders of magnitude).

The key to the error control strategy lies in how these tolerances are applied. For intermediate calculations within the PIMPLE loop, the solver can terminate using either the absolute or relative tolerance, whichever is met first. This provides computational efficiency by preventing the solver from over-working on a variable that is already well-converged. However, for the final correction step of each time step, the `relTol` is set to 0.

This deactivates the relative tolerance criterion, forcing the solver to adhere strictly to the absolute tolerance of $1e-9$. This "final check" ensures that, regardless of the intermediate steps, the solution written to disk at the end of every time step has met a consistent and stringent convergence standard. This robust strategy guarantees that the iteration error is kept to a negligible level throughout the simulation.

Control of Discretisation Error

Discretisation error is defined as the difference between the exact solution of the continuous governing equations and the solution of the algebraic system produced by the numerical method (Ferziger et al., 2020). This error is directly influenced by the quality of the computational mesh and the accuracy of the numerical schemes.

Ideally, a formal grid convergence study, involving simulations on at least three systematically refined meshes, would be conducted to quantify the magnitude of this error. However, due to the significant computational and time cost associated with generating and simulating multiple large-scale meshes for each valve lift configuration, such a study was beyond the scope of this thesis. In its absence, a strategy of proactive quality control was adopted to minimise discretisation error by ensuring a single, high-quality mesh was used in conjunction with robust and accurate numerical schemes.

The quality of the computational mesh was ensured through rigorous evaluation using OpenFOAM's `checkMesh` utility, as detailed previously. All meshes contained 2.2–3.2 million predominantly hexahedral cells (~80%), with non-orthogonality below 65° and negligible high-skewness cells (<0.0003%). These metrics, along with the overall `checkMesh` "OK" status, confirm that the meshes meet established quality standards and are suitable for accurate CFD simulations as previously detailed in Tables 7 and 8.

On top of mesh quality checks, the numerical schemes, defined in "system/fvSchemes", were chosen to balance accuracy with stability. Second-order `limitedLinear` schemes were used for key convective terms, which blend between high-accuracy

central-differencing and stable upwind schemes in regions of high gradients. This "limiting" action is crucial for preventing non-physical oscillations and ensuring the simulation remains robust.

Temporal discretisation error was controlled by using an adaptive time-stepping strategy, as defined in "system/controlDict." The solver adjusted the time step Δt dynamically to maintain a maximum Courant Number of 5. The Courant number is a dimensionless measure of how far fluid travels across a cell in one time step and keeping it below a stable limit is critical for transient accuracy. This adaptive approach ensures that the time step is always small enough to resolve the flow's evolution without compromising stability.

Assessment of Modelling Error

Having controlled iteration and discretisation errors to a satisfactory extent, the final step in the V&V hierarchy is the assessment of modelling error. This error is defined as the difference between the real-world flow physics and the exact solution of the chosen mathematical model (Ferziger et al., 2020). The final comparison between the simulation results and the experimental data in Chapter 4 is, therefore, fundamentally an evaluation of the magnitude of this modelling error. The primary sources of modelling error in this work are the RANS turbulence models.

As detailed in Section 2.4, RANS approach, while computationally necessary for industrial-scale problems, relies on models to close the governing equations. These models, including the $k - \epsilon$ and $k - \omega$ families, are not universal laws but are calibrated approximations, each with its own strengths and weaknesses. The choice of turbulence model can significantly impact the predicted flow field, particularly in complex flows like those in engine ports, which are characterised by: strong streamline curvature and swirl, large regions of flow separation and re-attachment (e.g., downstream of the valve seat).

To directly quantify the uncertainty introduced by this modelling choice, an extensive turbulence model sensitivity analysis was conducted as primarily presented in Sections 4.2.1. and 4.3. Rather than relying on a single model, the simulation for a representative case was performed with several widely used two-equation RANS models, including: Standard $k - \epsilon$ ($k - \epsilon$) , *Realizable $k - \epsilon$* , $k - \omega$ SST , $k - \omega$ SST, and $k - \omega$ (2006).

This comparative study serves a critical purpose: by running the exact same case setup with different turbulence models, the resulting spread in the key output quantities (such as mass flow rate) provides a quantitative measure of the modelling uncertainty. It establishes a confidence interval for the CFD predictions that is attributable solely to the turbulence closure.

In the two families of turbulence models, the following modelling assumption made in the near-wall regions, i.e., `epsilonWallFunction` for the $k - \epsilon$ family, and corresponding `omegaWallFunction` for the $k - \omega$ models. This approach avoids resolving the boundary layer all the way to the wall with an extremely fine mesh. Instead, it assumes that the flow in this region follows the universal "law of the wall." This is a necessary and standard engineering practice that makes the simulation computationally feasible, but it introduces a modelling error that is common to all the RANS models tested.

In conclusion, by following a structured V&V process, iteration and discretisation errors were controlled. The assessment of modelling error was then elevated from a simple description to a direct investigation through a comprehensive turbulence model sensitivity study. The final discrepancy observed between the CFD results and the experimental data can therefore be interpreted with a clear understanding of the uncertainty bounds introduced by the RANS modelling approach itself; Sections 4.2.1 and 4.3 have been dedicated to the analysis and quantification of this error.

3.3 Development of PyFlowBench Package

One of the key challenges identified in previous numerical Flow bench campaigns was the significant time investment required for each CFD simulation. Running a single case could take up to one or two weeks, making it impractical to execute comprehensive validation campaigns against large experimental datasets. This time constraint also hindered the ability to perform efficient DoE iterations, which are essential for rapid prototyping and optimisation of next-generation combustion system components. To address this bottleneck, the PyFlowBench Python package was developed as part of this thesis work. The primary objective of PyFlowBench is to automate and streamline the CFD simulation workflow, thereby enabling scalable, reproducible, and time-efficient simulation campaigns. The package encapsulates several core functionalities:

- **Automated mesh generation:** bash scripts that allows quick mesh generation with SHM.
- **Data transfer automation:** automated data transfer from the experimental dataset into the numerical Flow bench custom dictionary “FlowBenchDict” necessary for configuration selection, valve lifts (cases), boundary conditions such as, inlet & outlet pressure, temperature, experimental mass flow rate reference values from GKN and Wärtsilä sensors and the values from in-cylinder temperature sensor.
- **Preprocessing automation:** The tool performs a series of pre-run validation checks before initiating any simulation. It then automates the entire case setup process, including mesh handling for both continuous volume meshes (SnappyHexMesh) and modular structured meshes (GridPro). Directory management is handled automatically, and the package also supports mesh mirroring, mesh merging, and patch renaming where applicable.
- **Simulation submission:** integrates with the HPC environment to manage job submission/execution.
- **Postprocessing routines:** calculating and plotting mass flow rates, temperature, and heatmaps for analysis.

- **Cleanup and memory management:** post analysis, another bash script reconstructs the field data and cleans up all the unnecessary directories and files.

By significantly reducing manual intervention, PyFlowBench not only improves efficiency but also ensures consistency across simulation cases. It facilitates large-scale simulation campaigns needed for rigorous model validation and supports iterative development workflows, such as sensitivity studies and machine learning-based surrogate model training for a future work.

In summary, PyFlowBench serves as an essential enabler for the numerical Flow bench methodology, bridging the gap between computational demands and practical engineering timelines. Figure 33 shows the workflow of the package.

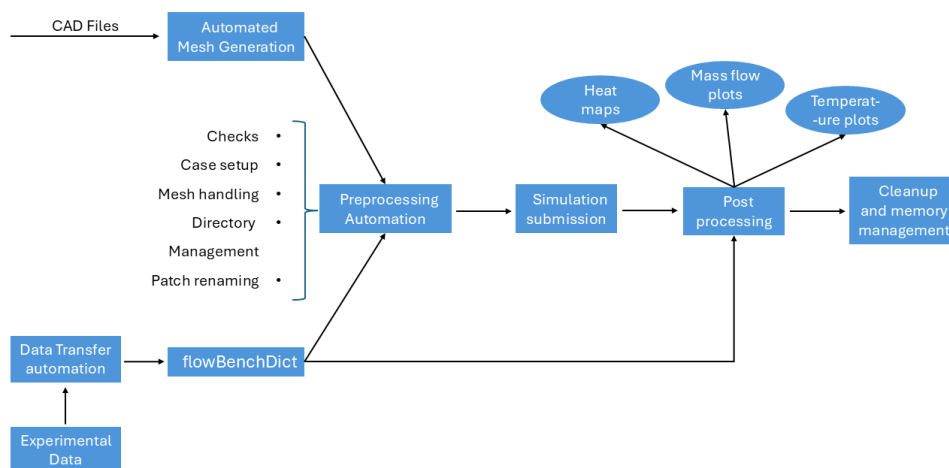


Figure 33. Flow chart of PyFlowBench workflow.

Importance of PyFlowBench

The PyFlowBench package plays a pivotal role in enhancing the efficiency, scalability, and reliability of the numerical Flow bench workflow for the Wärtsilä engine component simulations.

Previously, mesh generation was carried out using GridPro, a software known for producing high-quality meshes but requiring extensive manual effort and time. While GridPro meshes remain valuable for their geometric fidelity, PyFlowBench introduces a scripted meshing workflow using snappyHexMesh, which significantly accelerates the process. Though SHM typically yields lower-quality meshes compared to GridPro, its integration into the automated pipeline allows for rapid case setup, making it highly suitable for DoE and iterative validation runs.

Another time-saving feature is the automated data transfer module, which populates the simulation dictionary (`FlowBenchDict`) directly from experimental datasets. This reduces the labour involved in manual data entry, especially when multiple cases are run in parallel, and minimises the risk of human error.

The `preprocess.py` script handles both preprocessing and simulation submission. It performs essential validation checks such as confirming the presence and correctness of boundary conditions and sourcing the appropriate OpenFOAM environment. It also clones a predefined case structure (`FlowbenchBase`) into each simulation directory, ensuring consistent case configuration across varied mesh inputs. Moreover, it supports parallelisation of simulations across a specified number of HPC cores, enabling efficient batch processing and large-scale case management.

Finally, postprocessing scripts automate the extraction of critical performance metrics such as mass flow rate and temperature at the cylinder head. These can be executed with a single command without any additional input from the user since the scripts read the necessary data from previously populated `FlowBenchDict` accelerating the analysis phase while maintaining result reproducibility.

In conclusion, the PyFlowBench package has proven to be a critical enabler in achieving automation, consistency, and scalability across the entire CFD workflow. From mesh generation and boundary condition assignment to simulation execution, result

postprocessing, and cleanup, each component of the package contributes to reducing manual effort and ensuring repeatability. Moreover, PyFlowBench lays the groundwork for high-throughput design studies supporting rapid prototyping and advanced analysis in future research and industrial applications.

4 Results and Discussion

This chapter presents the key results obtained through the CFD simulations of the numerical FlowBench framework. The primary focus is on:

- evaluating the performance of the CFD simulations
- validating them against experimental data
- identifying the sources of uncertainty in both numerical and experimental flow bench.

It should be noted that default coefficients for all the numerical simulation turbulence models have been used.

The results are structured into distinct subsections corresponding to different stages of the study starting with the notes on experimental uncertainties, turbulence model sensitivity analysis, effects of pressure ratio on mass flow rate (\dot{m}), effects of valve lift on \dot{m} , validation of numerical temperature against the experimental values and 3D flow characteristics of different CHs and valve openings.

4.1 Initial Validation and Emerging Experimental Uncertainties

The initial validation phase of the CFD model, intended to confirm numerical accuracy, served a more critical diagnostic purpose: it systematically revealed significant uncertainties within the experimental dataset that were not immediately apparent. The initial simulations were conducted for CH1 Pressure Ratio of 1.1, exhaust configuration, which yielded encouraging results. As shown in Figure 34, deviations from experimental data were generally contained within a $\pm 5\%$ margin combining the performance of different turbulence models. Similar consistency was observed for PR=1.5 under varying absolute pressure conditions shown in Figure 35, providing initial confidence in the numerical setup.




Figure 34. Lifts from minimally open to fully open, exhaust flow pressure ratio of 1.1 – CH1.

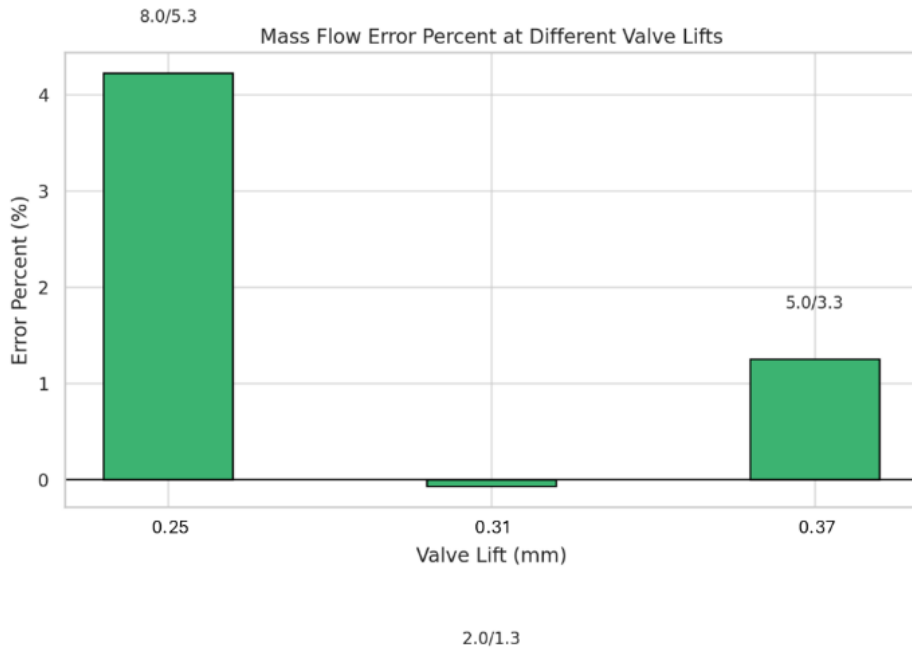


Figure 35. CH1 exhaust flow configuration at normalised lifts of 0.25, 0.31 and 0.37 with pressure ratio of 1.5 and varying the absolute Pin/Pout.

Encouraged by the promising exhaust results, validation of the intake flow cases was undertaken. A closer agreement was anticipated, given that intake flows are characterised by cooler, less complex jet structures compared to the thin, wall-hugging jets of the exhaust. Contrary to expectations, the intake cases showed significantly greater divergence from the experimental values.

This discrepancy, particularly on the intake side, prompted a thorough review of the numerical setup. All simulation configurations were checked and refined, but the deviations persisted. This finding strongly suggested that the error source was not primarily numerical, but was instead inherent to the experimental data being used for comparison.

A critical review of the raw experimental data revealed direct evidence of measurement discrepancies. These inconsistencies point toward the presence of significant systematic errors (bias), a class of error that, as discussed previously, can originate from instrumental or installation effects (Figiola & Beasley, 2020).

A definitive example of this is the direct comparison of the Wärtsilä and GKN sensors measuring the same intake flow condition for the normalised lift of 1.0 and pressure ratio of 1.1. The mass flow rate reported by Wärtsilä's sensor was 1.5402010679, while the GKN sensor reported 1.3481249809. This constitutes a relative deviation of 14.2%, calculated as:

$$Percent_{error} = \frac{Wärtsilä \text{ mass flow} - GKN \text{ mass flow}}{GKN \text{ mass flow}} * 100 \approx 14.2477 \%. \quad (28)$$

A discrepancy of this magnitude, illustrated in Figure 36, between two independent instruments operating within their valid ranges cannot be attributed to random statistical scatter. It is a clear indication of a large systematic bias between the two measurement systems. Figure 36 showed that the uncertainty from using different experimental reference data, i.e., Wärtsilä vs GKN, had a greater impact on the calculated error than the choice of turbulence model. This finding was pivotal, as it demonstrated that the

experimental data contained uncertainties far greater than those of the numerical model it was meant to validate. It is noteworthy that these early observations are not conclusive, particularly regarding the performance of different turbulence models. However, they laid the groundwork to uncover potential experimental discrepancies. The final results from turbulence sensitivity analysis are discussed in Sections 4.2.1 and 4.3. Therefore, the following discussion uses these initial findings as a motivation to explore the experimental uncertainties that were subsequently investigated starting with more in-depth assessment of intake flow simulations.




Figure 36. Initial turbulence model sensitivity analysis for intake flow CH1 at normalised lift of 1.0. The figure demonstrates that the choice of experimental reference data (Wärtsilä vs. GKN) introduces a larger variation in the calculated error than the choice of turbulence model. This highlights the dominant role of experimental uncertainty in the validation process.

4.2 Intake Flow Configuration

The intake flow simulations were approached systematically, with the following objectives:

- conducting a sensitivity analysis of various RANS turbulence models with respect to predicted mass flow rate (\dot{m})
- examining the effect of valve lift on \dot{m}
- evaluating the influence of pressure ratio on \dot{m}
- performing a comprehensive valve lift sweep
- validating the predicted temperature at the numerical probe location against the experimental temperature sensor readings.

4.2.1 RANS Turbulence Models Sensitivity Analysis on Mass Flow Rate

During the turbulence model sensitivity analysis, five RANS turbulence models were selected and evaluated under two representative valve lifts and two distinct pressure ratios as observed in Figures 37 and 38. Upon completing the simulations and analysing the results, several key observations were made:

- For a given valve lift, better agreement with experiment was consistently observed at the higher pressure ratio of 1.5 across nearly all turbulence models.
- The simulation results showed closer alignment with the GKN sensor compared to those from Wärtsilä's.
- $k - \varepsilon$ turbulence model seems promising with everything considered.
- An improvement in model accuracy was observed with increasing valve lift, indicating that higher lifts tend to yield more reliable comparisons between CFD predictions and experimental data.




Figure 37. Five turbulence models performance against GKN Sensors at two lifts and pressure ratios. It can be observed that accuracy improved at higher lift and pressure ratio, the $k - \epsilon$ model performed best overall, and simulations aligned more closely with GKN data than Wärtasilä's.




Figure 38. Five turbulence models performance against Wärtasilä sensors at two different lifts and pressure ratios.

Following the initial sensitivity analysis, the study was expanded to include a broader range of representative test points using the most promising turbulence models. This extension aimed to further validate the earlier observations, namely, the tendency for better agreement with GKN sensor data, improved accuracy at higher pressure ratios, more consistent performance at higher valve lifts and $k - \varepsilon$ potential as the final candidate out of all the turbulence models. Specifically, additional cases at a small, normalised lift of 0.14 and full lift of 1.0 were simulated under pressure ratios of 1.1 and 1.5 using the $k\omega SST$, $k - \varepsilon$ and $RNG K - \varepsilon$ turbulence models.

As illustrated in Figure 39, a higher pressure ratio of 1.5, consistently led to improved agreement across all valve lifts. Moreover, the earlier observation that CFD predictions aligned more closely with GKN sensor data than with Wärtsilä's measurements held true throughout this extended analysis. While the trend of improved accuracy with increasing valve lift was not conclusively evident in this subset of cases, it is clearly demonstrated in the subsequent "valve sweep analysis" observed in Figure 43, where the pattern becomes more pronounced.

The expanded investigation, also, conclusively demonstrated the overall higher performance of the standard $k - \varepsilon$ model relative to its counterparts. Although an exception was noted at the normalised lift of 0.57, the error percentage is sufficiently small (-1.15%), that it makes no practical difference whether $k\omega SST$ or $RNG K\varepsilon$ surpass it by a small margin. Based on these results, the $k - \varepsilon$ turbulence model was considered the optimal modelling choice for the pressure ratio and valve lift sweep investigations and was therefore utilised for all subsequent validation efforts.




Figure 39. The extended comparison of 3 turbulence models against 2 sensors at 4 distinct lifts and 2 pressure ratios. The empty rows in Wärtsilä error heat map are cases in which this sensor is not valid. It can be concluded that GKN data aligns more closely with CFD, higher pressure ratio improves agreement, and the standard $k - \epsilon$ model consistently outperforms the others.

Also, regarding the temperature validation, a reverse trend was observed to that of mass flow rate, i.e., with increasing pressure ratio, higher deviation was observed. This deviation can be due to no heat flux across the wall with fixing adiabatic assumption in the boundary conditions. However, $k - \epsilon$ still outperforms its peers, thus, indicating a higher performance overall as observed in Figure 40. A potential explanation behind this observation can be found in section 4.5.4.

Temp Validation for Small vs Full Lifts across different Turb. models




Figure 40. Temperature validation against the experimental temperature sensor shows higher deviations at higher pressure ratio and $k - \epsilon$ outperforming its counterparts.

4.2.2 Effects of Pressure Ratio on Mass Flow Rate

Building on earlier observations that higher pressure ratios improved agreement in mass flow rate predictions but increased deviations in temperature, a targeted study was carried out to examine this trade-off more systematically. To this end, a dedicated pressure ratio sweep was performed at pressure ratios of 1.1, 1.5, 2.0, 3.0, and 4.0. The results, presented in Figure 41, reveal several key insights:

- The earlier trend was reaffirmed; Numerical predictions of mass flow rate aligned more closely with the GKN sensor measurements than with those of Wärtsilä's instrumentation.
- As pressure ratio increased, the accuracy of the predicted mass flow rate improved. This was reflected in a systematic decrease in the percentage error between simulation and experimental data.
- Conversely, an increase in pressure ratio was associated with a growing discrepancy in temperature predictions, as illustrated in Figure 42. The temperature

error percentage showed a clear upward trend with increasing pressure ratio, indicating limitations in thermal accuracy under high-pressure conditions.

- At the expected (exp) pressure ratio of 4.0, the actual (act) pressure ratio achieved in the experiment failed to go any higher. Instead, it plateaued around 2.3 for all higher pressure ratios. This was the result of high mass flow causing back pressure at the waste gate, which obstructed the outflow. The pressure at the outlet increased significantly consequently, reaching 3.74 bar at inlet pressure (Pin) of 8.0 bar, while it was supposed to be maintained at 1.0 bar. This limited any higher pressure ratio than 2.3.

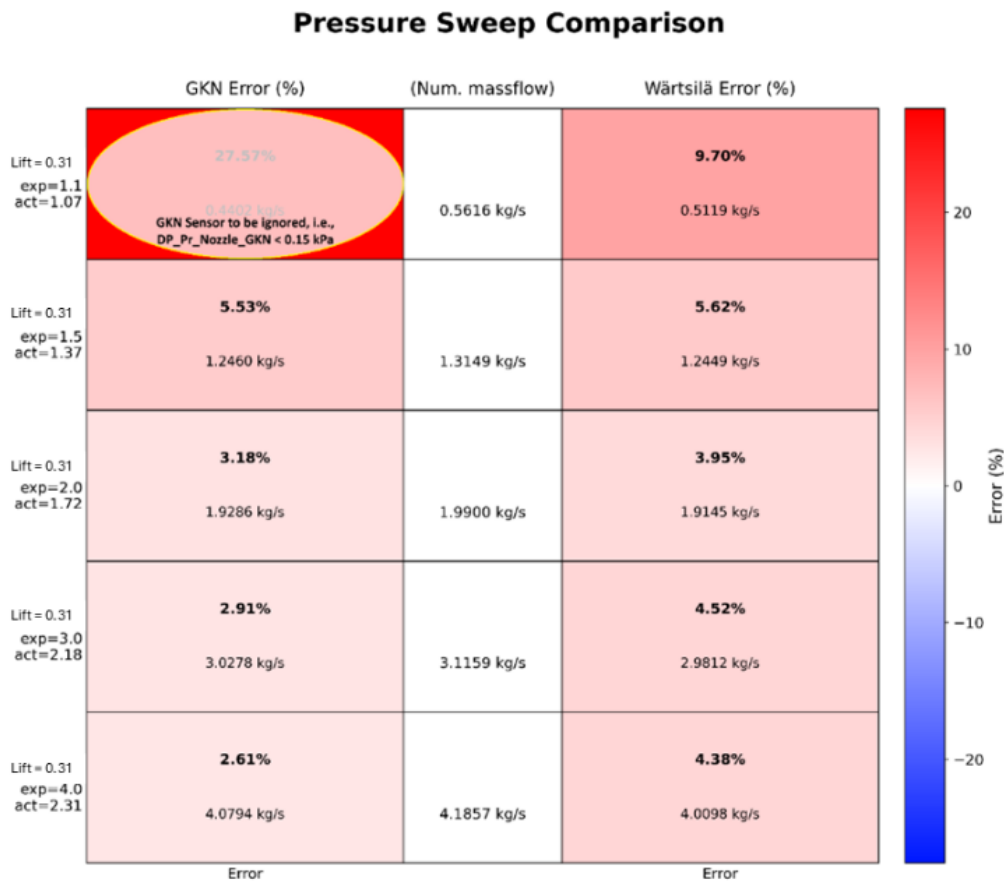


Figure 41. Mass flow error during a pressure ratio sweep (exp. Pressure ratio 1.1–4.0). CFD matched GKN sensors better than Wärtsilä's. Accuracy improved with increasing pressure ratio. Actual pressure ratio plateaued at ~2.3 due to outlet backpressure.




Figure 42. Temperature error across the same pressure ratio sweep. Temperature predictions worsened with higher pressure ratio, indicating higher thermal inaccuracies under high-pressure conditions.

4.2.3 Effects of Valve Lifts on Mass Flow Rate

Figure 43 illustrates the effects of valve lift on the predicted mass flow rate at the pressure ratio of 1.5. A key observation was the increasing convergence between the numerical results and the experimental data from the GKN sensor at higher valve lifts. While a noticeable deviation existed at low lifts, the two datasets showed excellent agreement as the valve lift increased. It is further exemplified in Figure 44, where the mass flow rate has been normalised over its lift. The line graph shows a fairly linear decline excluding the anomalies. The anomalies occur at 42.8 % and 45.5 % lift, where the two sensors already disagree most strongly (-10.19 % and -8.93 %, respectively).

Therefore, the following two conclusions are drawn:

1. With increasing lifts and mass flow rate, there is less experimental uncertainty, thus, leading to closer agreement between the numerical and measured mass-flow rates. It is explained in detail in Section 4.5.3.

2. Figures 45 and 46 confirm that higher absolute inlet pressures also improve the match, even when the pressure ratio itself remains unchanged as the error percentage is bound to under $\pm 5\%$.

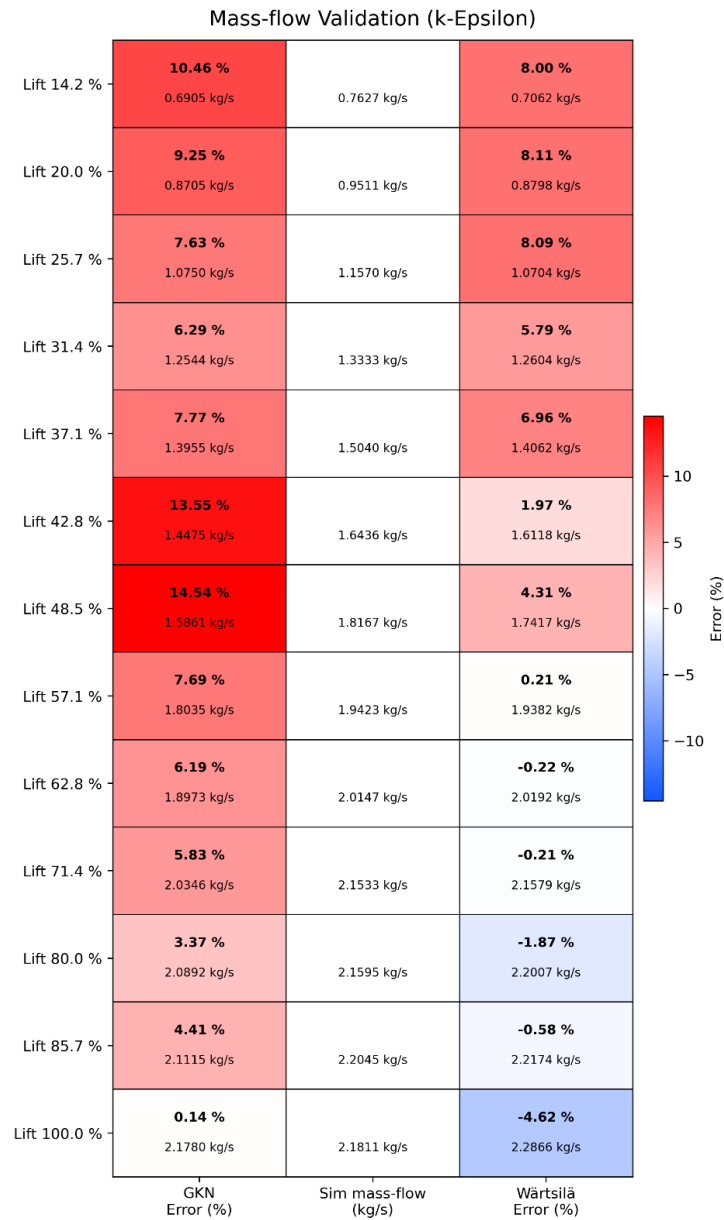


Figure 43. Valve lift sweep from minimally open to 100% open at pressure ratio of 1.5 with P_{in} 1.5 bar and P_{out} set to ambient pressure. The main takeaway is that with increasing lift, which corresponds to higher \dot{m} , there is better matching between experimental and numerical values.




Figure 44. Valve lift sweep from minimally open to 100% open at pressure ratio of 1.5 with inlet pressure (P_{in}) 1.5 and outlet pressure (P_{out}) set to ambient pressure showing mass flow rate (\dot{m}) normalised over lift vs their respective error percentage. Ignoring the anomalies encircled in yellow, there is a rather linear decline in the error percentage as the valve lift grows.




Figure 45. Valve lift sweep from minimally open to 100% open at pressure ratio of 1.5 with Pin 8.0 and Pout 5.3 bar. At higher absolute inlet pressure, there is lower uncertainty with the sensors, therefore, the deviation between the CFD and experimental values are small.

This trend does not hold under all operating conditions. As shown in Figure 45, at a higher inlet pressure, i.e., Pin = 8.0 bar, the resulting increase in mass flow rate alters the

system's behaviour. The error no longer decreases with lift; instead, it remains relatively low and stable across the entire lift range which is consistently below 4.6%.

This stabilised error can be attributed to the improved accuracy of the experimental sensors under high-flow conditions since it operates within its intended flow regimes. This point is further explained in Section 4.5.1. Therefore, in these cases the predictive accuracy of the CFD model becomes largely independent of valve lift and the numerical predictions stabilise in terms of accuracy. The error percentage oscillation with respect to lift can be observed in Figure 46.




Figure 46. Valve lift sweep from minimally open to 100% open at pressure ratio of 1.5 with Pin 8.0 and Pout 5.3 bar showing mass flow rate (\dot{m}) normalised over the lift vs its respective error percentage. The error percentage oscillates and there is no clear decrease w.r.t lift.

Regarding the temperature validation of the valve sweep cases, previously observed trends remain consistent: as the pressure ratio increases, leading to higher mass flow

due to elevated inlet pressure, the temperature deviation also increases. This is clearly seen by comparing Figures 47 and 48. At lower absolute pressure (e.g., $P_{in} = 1.5$ bar in Figure 47), the temperature error percentage is relatively small, whereas higher inlet pressure ($P_{in} = 8$ bar in Figure 48) results in comparatively larger discrepancies. Further explanation is provided in section 4.5.4.




Figure 47. Temperature validation for the valve sweep at pressure ratio of 1.5 with P_{in} 1.5 bar and P_{out} set to ambient pressure. Lower absolute inlet pressure results in smaller temperature errors.

Figure 47 shows no pronounced upward trend in numerical temperature as the valve lift increases. Temperature deviation rises until it reaches lift that is 57 % open, then it marginally reduces towards full opening. This behaviour indicates that although temperature deviation also increases with the valve lift, the absolute inlet pressure, perhaps, has a stronger influence on temperature variation than the valve lift. However, as expected, the highest deviation is observed when there are the most uncertainties on the experimental side between the two sensors at mid lifts.

Cylinder-Head Temperature Validation (k-Epsilon)

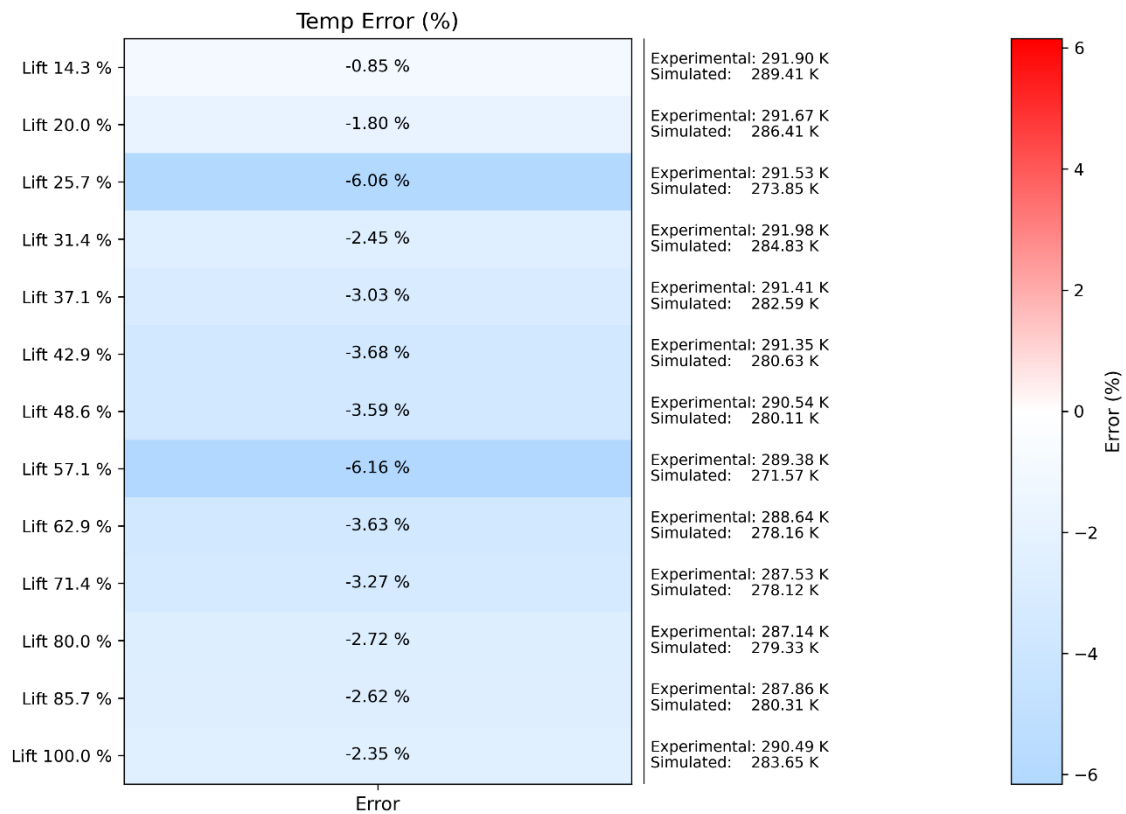


Figure 48. Temperature validation for the valve sweep at pressure ratio of 1.5 with Pin 8.0 bar and Pout 5.3 bar. Higher absolute inlet pressure leads to increased temperature deviation, especially at mid-range lifts.

Figure 48 follows a similar trend to that observed in Figure 46. It shows that although the error percentage in temperature increases with valve lift, the effect is less pronounced than the variation caused by changes in the absolute pressure. In other words, while lift does contribute to increasing thermal deviation, its impact is not as dominant as that of absolute inlet pressure as indicated in Figure 42 during pressure ratio sweep analysis. Nevertheless, higher deviations are evident in Figure 48, as expected, the elevated mass flow rate resulting from higher absolute inlet pressure leads to a greater amount of mass being transported into the domain. This increases interaction with the walls, making accurate wall modelling more critical. In the current numerical setup, adiabatic wall conditions (i.e., zeroGradient temperature) are assumed, meaning no heat flux or energy exchange occurs between the walls and the air.

In reality, however, the walls are likely to remain near ambient room temperature with some spatial variation across the cross-section. This realistic wall temperature would cause thermal energy to be transferred via heat flux from the walls to the air, which cools due to rapid expansion and the associated pressure drop after passing through the valves. This crucial heat exchange phenomenon was not accounted for in the present model.

4.3 Exhaust Flow Configuration

With the insights gained from the intake flow analysis, the exhaust flow configuration was revisited to re-evaluate selected test points. Recognising that greater uncertainty exists at lower pressure ratios, a targeted analysis was conducted at pressure ratio of 1.5 using three turbulence models across three representative valve lifts. As shown in Figure 49, the deviation between numerical and experimental mass flow rates remained below 4.21% across all cases. At lower valve lifts, the *RNG* $k - \varepsilon$ model exhibited the best performance and consistent with the trends observed in the intake flow configuration, the standard $k - \varepsilon$ model demonstrated a reduction in error percentage with increasing valve lift.

Regarding temperature validation, Figure 51 illustrates that temperature deviation increases with valve lift, an observation that further supports the previously identified trend in intake flow simulations.




Figure 49. Error percentage and mass flow rate of exhaust flow simulations at pressure ratio 1.5 with 3 turbulence models at 3 representative lifts measured against GKN’s sensor. While all the turbulence models showed acceptable performance, *RNG k – ε* outperformed at lower lifts and the error percentage for *k – ε* decreased with increasing lifts.




Figure 50. Error percentage and mass flow rate of exhaust flow simulations at pressure ratio 1.5 with 3 turbulence models at 3 representative lifts measured against Wärtsilä’s sensor. Overall, this sensor underperforms compared to GKN’s.




Figure 51. Temperature validation for exhaust flow at pressure ratio 1.5 across 3 lifts and turbulence models. With increasing lifts, increasing temperature deviation is observed.

Due to time constraints, no further systematic analysis could be performed on exhaust flow configurations.

4.4 Comparison of Different Cylinder Head Designs

One of the key advantages of three-dimensional numerical simulations is the ability to visualise the detailed flow field in high resolution. These simulations provide instantaneous access to all flow variables solved during the computation, offering insights that are often difficult or impossible to obtain through experiments alone.

In the numerical flow bench simulations, variables such as turbulent kinetic energy (k), dissipation rate (ε), specific dissipation rate (ω), pressure (p), temperature (T), and velocity (U) were computed and stored. These variables can be post-processed and visualised to uncover complex flow behaviours within the domain.

In particular, the 3D visualisations enable meaningful comparisons between different cylinder head configurations and valve arrangements. They reveal asymmetries, flow separations, recirculation zones, and mixing behaviour that may otherwise remain hidden.

The following sections present the following comparative analyses:

- CH1 vs. CH3: Intake and exhaust configurations with all valves open.
- CH1 Intake Setup: Comparison between cases with only valve A open vs. only valve B open.
- CH1 Exhaust Setup: Comparison between cases with only valve C open vs. only valve D open.

4.4.1 Cylinder Head 3 vs. Cylinder Head 1

To recapitulate, CH3 is a five-valve cylinder head featuring three intake valves and two exhaust valves. In contrast, CH1 is a four-valve configuration with two intake and two exhaust valves. Beyond the valve count, the two cylinder heads also differ in valve size, seat angle, and geometrical design, all of which influence the flow characteristics and overall performance.

Intake Flow Comparison (CH3 vs. CH1) at Normalised Lift of 0.25

At the normalised valve lift of 0.257, the GKN sensor recorded a normalised mass flow rate of 0.703 for CH3 and 0.465 for CH1. CH3's higher mass flow is evident from the broader envelope of high-velocity regions (peaking $\approx 290 \text{ m s}^{-1}$) around the valve seats observable in Figure 52. Upstream of the valves the riser-pipe patterns are broadly similar for both the CHs. In CH3, the interaction between the centre valve and its adjacent valves produces the highest velocities within the cylinder and keep its momentum longer. A similar interaction pattern can be seen between the intake valves in CH1. Overall, CH3 generates a higher tumble ratio.




Figure 52. Side by side comparison of CH3 (left) vs CH1(right). The first two are the front and the bottom two are the back view. The high velocity region around the valve seat is more dominant in CH3 compared to CH1. Also, riser-pipe patterns are broadly similar for both and intake valve-intake valve interactions lead to higher velocity inside the cylinder.

Exhaust Flow Comparison (CH3 vs. CH1) at Normalised Lift of 0.714

Figure 55 illustrates the exhaust flow for CH3 versus CH1 at normalised lift of 0.714. CH1 exhibits significantly higher flow velocity at the valve seat area. However, the interaction between CH3 exhaust valves results in a sustained, higher-velocity flow further upstream past the valves. Furthermore, the sharp curvature of the exhaust port induces recirculation zones just before the bend, which could impact flow efficiency.

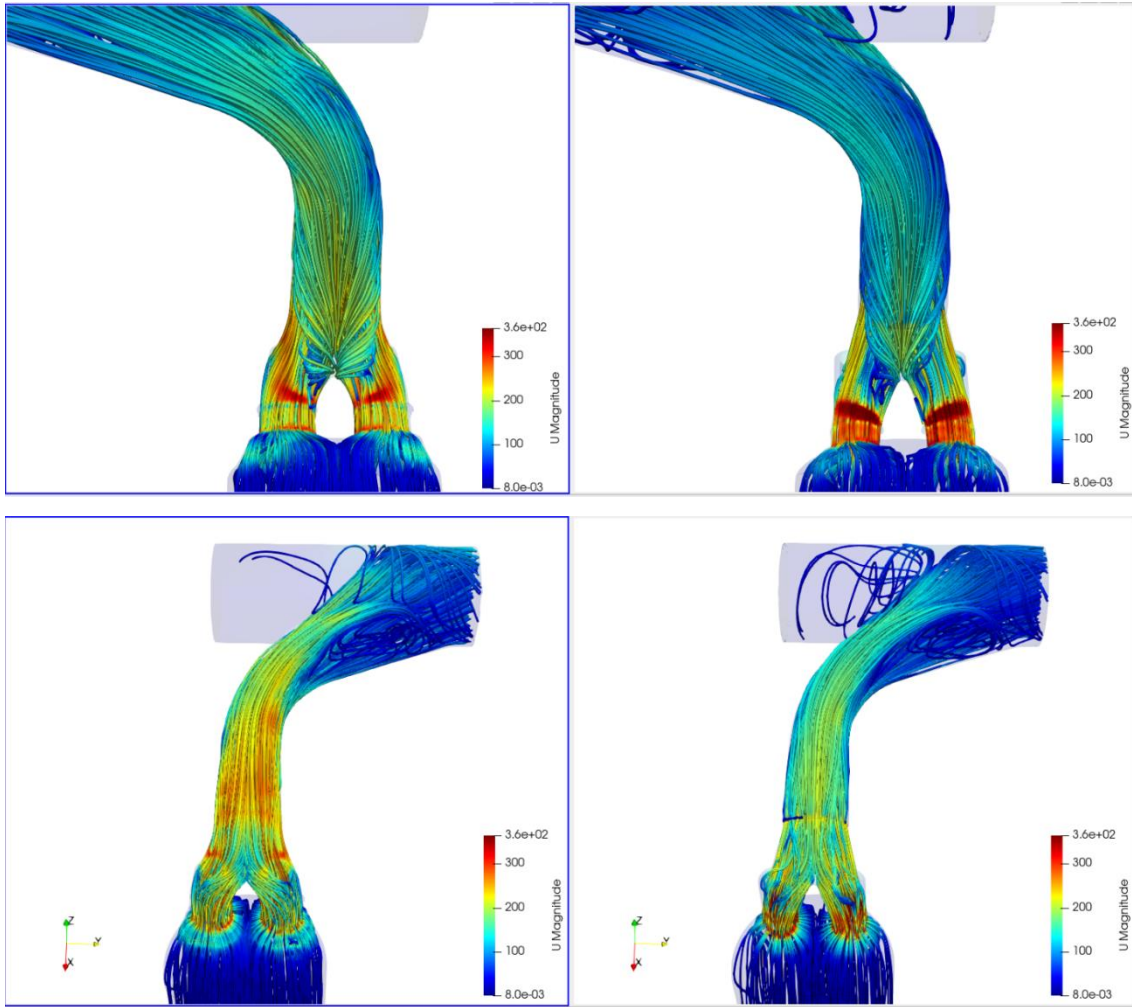


Figure 53. Exhaust flow comparison of CH3 (left) vs CH1 (right) at normalised lift of 0.714. Although CH1 shows increased flow velocity at the valve seat area, CH3's higher velocity is sustained longer in the exhaust port. Also, a re-circulation zone can be observed right before the bend.

4.4.2 Cylinder Head 1 Individual Intake Valves and Exhaust Valves

CH1 Intake: Static Valve A vs. Static Valve B at normalised lift of 0.25

Figure 53 depicts CH1 operating with a single intake valve at normalised lift of 0.257. When only valve A is open (left frame) the jet accelerates to $\approx 300 \text{ m s}^{-1}$ at the seat, leaves the port almost tangentially toward the +x wall, and adheres to the liner through roughly half a revolution before spiralling downward. The resulting flow field is

dominated by a coherent clockwise swirl (viewed from +z); tumble is negligible and a low-velocity recirculation pocket forms beneath the closed B valve on the opposite side. Opening only valve B (right frame) produces the exact mirror image: the jet attaches to the $-x$ wall, rotates anti-clockwise, and shifts the low-momentum region under the closed A valve. The mass flow rate differs by 2.3 % and the peak velocity is exactly the same, confirming that the two ports are aerodynamically balanced.

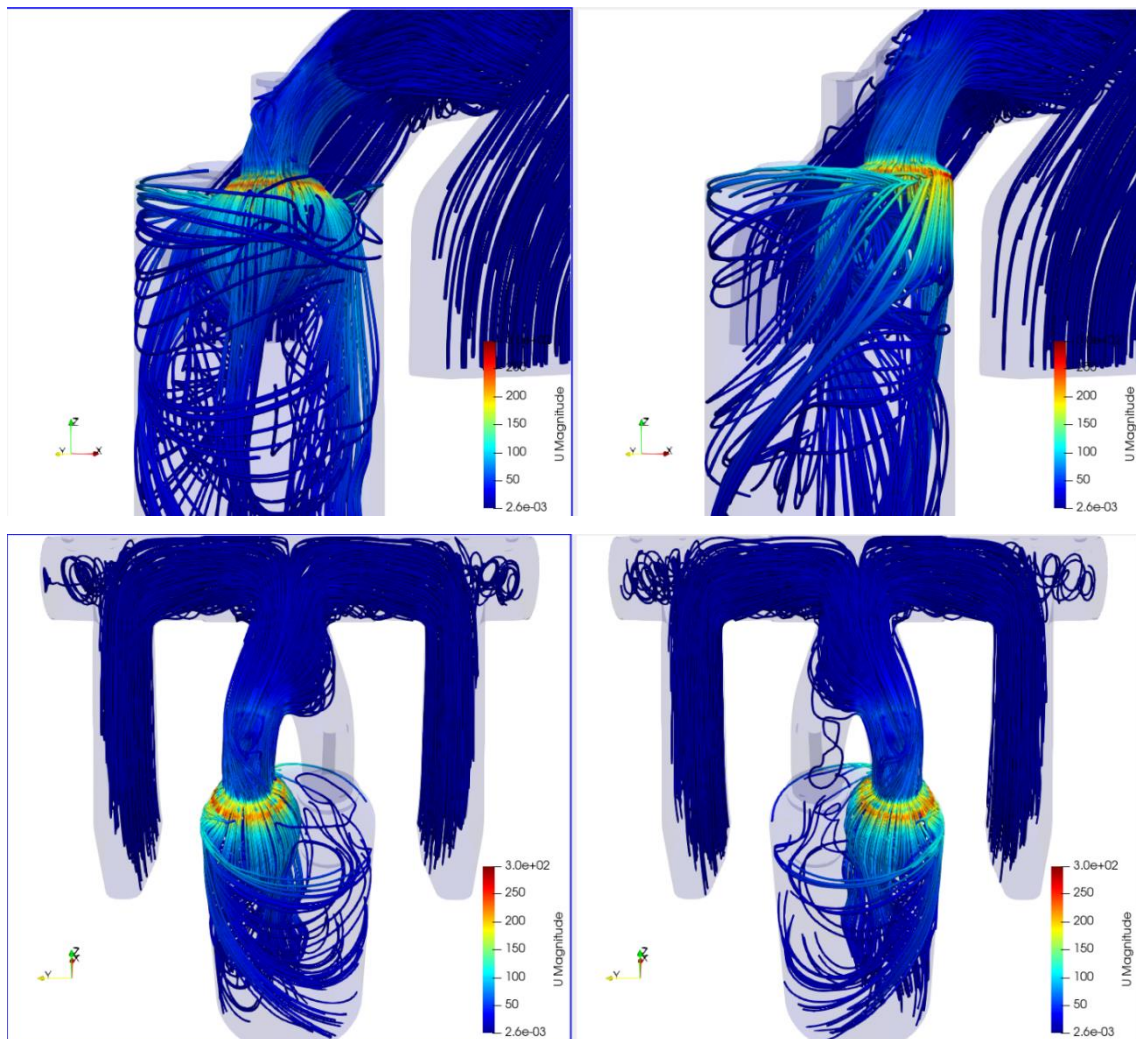


Figure 54. Comparison of flow structures in the CH1 intake for single-valve operation at a normalised lift of 0.257. Opening Valve A (left) versus Valve B (right) produces near-perfect mirror-image flow patterns dominated by a coherent swirl, demonstrating the port's symmetrical design.

CH1 Exhaust: Static Valve A vs. Static Valve B at normalised lift of 0.31

Figure 54 depicts CH1 operating with a single exhaust valve at the normalised lift of 0.314. There is a rather noticeable swirl in the flow exiting the open exhaust valve for both C and D. And a strong flow to the wall opposite to the open exhaust valve, and it continues upstream in the port.

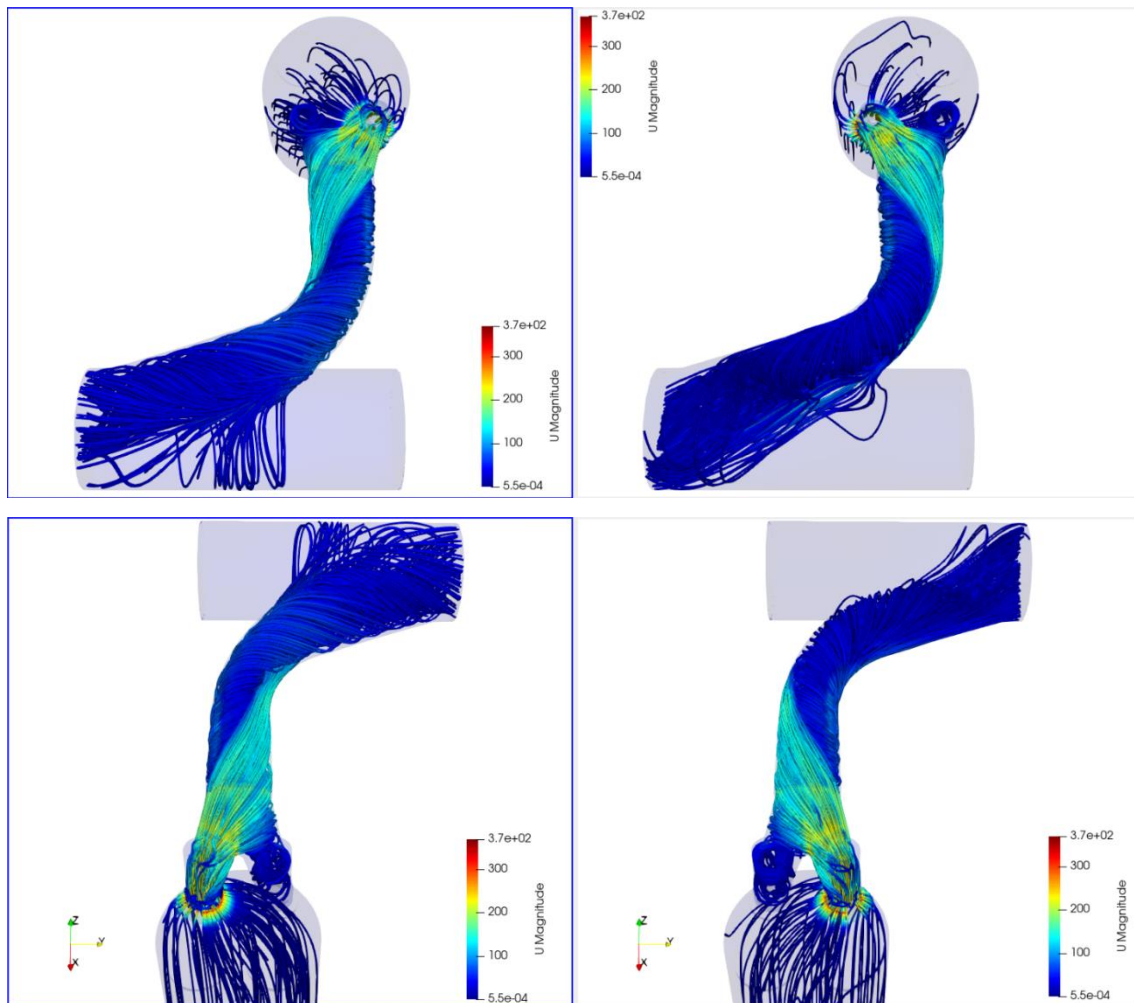


Figure 55. CH1 exhaust valve performance – Static valve C (left) against Static valve D (right). The flow shows a characteristic swirl at the valve exit and a strong, wall-impinging jet that travels upstream into the port.

4.5 Key Findings for Experimental and Numerical Discrepancies

This subsection consolidates the key observations from the validation campaign, linking them to their most plausible physical and instrumentation-related causes. Based on this analysis, specific areas for improvement in both the numerical model and the experimental setup are identified. Finally, concrete recommendations for future work are proposed to address these findings.

4.5.1 Influence of Absolute Pressure on Sensor Accuracy

Both the intake and exhaust test series begin at a pressure ratio of 1.1; however, the absolute inlet pressures differ, 1.5 bar for the intake and 2.0 bar for the exhaust. Under the ideal gas assumption, density increases with absolute pressure, and velocity is governed by the pressure gradient (ΔP across the domain). The higher absolute pressure in the exhaust configuration results in greater mass flow, which enhances the signal strength and reduces relative uncertainty in the measurements. This is clearly illustrated in Figure 56.

In contrast, the lower mass flow rate in the intake cases increases the noise-to-signal ratio, amplifying the impact of fixed absolute error inherent to differential-pressure (Bernoulli-type) sensors. These sensors infer velocity from the difference between total and static pressure ($P_{total} - P_{static} = P_{dynamic}$). At low velocities, this pressure difference becomes small, making the measurement more sensitive to noise and thus less reliable.

This effect is further evident at pressure ratio of 1.5, where two different combinations of absolute P_{in} and P_{out} were tested. The case with higher absolute pressures showed better agreement with CFD predictions, reinforcing the role of pressure level in improving measurement fidelity as shown in Figure 45.

It is clear that absolute inlet pressure is the defining factor for characterising the mass flow rate and consequently the uncertainty quantification of the flow sensors rather than pressure ratio. For instance, there can be multiple test points with the same pressure ratio but their measurements confidence would significantly vary depending on the absolute inlet pressure, cf. Section [4.2.3].




Figure 56. First batch of experimental test points for intake and exhaust configurations at a pressure ratio of 1.1. Intake tests were conducted with an inlet pressure of 1.5 bar and outlet pressure of 1.4 bar, while exhaust tests used 2.0 bar inlet and 1.8 bar outlet conditions.

4.5.2 Sensor Placement and Instrumentation Limitations

GKN Orifice Plate Sensor

The GKN orifice plate sensor may have been oversized for low pressure ratio conditions, as it is originally designed for high-throughput jet engine test rigs. Its performance improves significantly with increasing inlet pressure as the flow moves across the orifice. In particular, the sensor becomes unreliable when the differential pressure (ΔP_{Nozzle_GKN}) falls below 0.15 kPa, a condition that occurs at several valve lifts under pressure ratio of 1.1. This suggests that the sensor is not well-suited for capturing low-

flow scenarios in the current test configuration. Also, as noted in Section 3.1.3, the lack of a detailed calibration envelope for this sensor prevented a baseline uncertainty analysis, meaning its precise uncertainty range under these conditions remains unquantified.

Wärtsilä Pitot-Static Tube

Although it remains within its range at low flow rates, it is installed only $\sim 5 \times D$ downstream of a bend in the intake manifold (Figure 57), whereas as per the instrumentation and measurement expert and ITABAR-Sensor Type IBR Operating Manual (n.d.), it is recommended to have $\geq 7 \times D$ upstream and $\geq 3 \times D$ downstream. This suboptimal placement exposes the sensor to swirl and secondary flow structures, which are particularly prominent at low valve lifts. Additionally, the intake setup includes another bend in the vertical (y) direction, further compounding flow disturbances shown in Figure 58. In contrast, the exhaust configuration features a straight sensor placement, minimising such uncertainties shown in Figure 59. This suboptimal placement in the intake setup, combined with the lack of a specific calibration for these exact experimental conditions (as discussed in Section 3.1.3), means the true uncertainty of this measurement is likely higher than the theoretical baseline, especially at low valve lifts where flow distortions are most prominent.

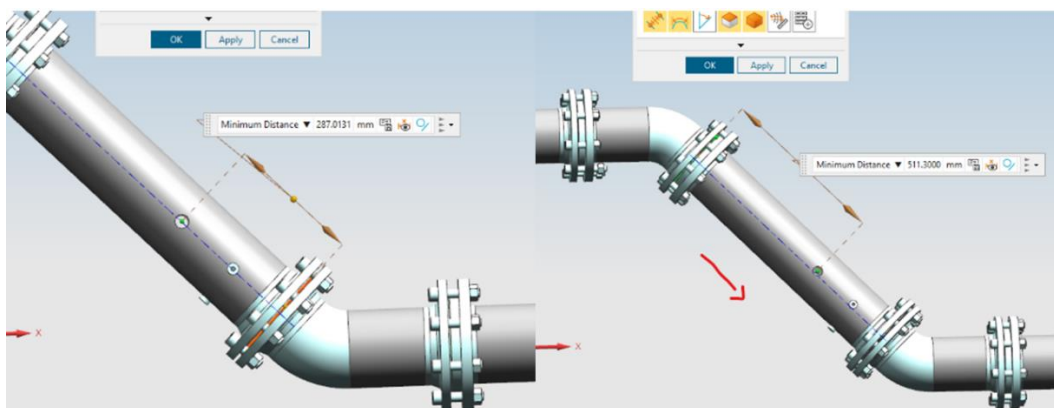


Figure 57. The suboptimal location of the sensor in Wärtsilä's inlet pipe.




Figure 58. Top view of the CH1 intake flow configuration, highlighting the bend in the y-direction where the Pitot-static tube is located. This additional bend introduces a potential source of experimental uncertainty. Proprietary components are intentionally blurred in accordance with Wärtsilä confidentiality requirements.




Figure 59. Top view of exhaust flow configuration for CH1. Wärtsilä's inlet pipe goes straight down under the bench. Proprietary components are intentionally blurred in accordance with Wärtsilä confidentiality requirements.

4.5.3 Effects of Absolute Pressure and Valve Lift on Mass-Flow Agreement

Higher pressure ratio, resulting from increased inlet pressure, achieves better agreement. Across all configurations, raising pressure ratio from 1.1 to 1.5, 2.0 and higher consistently reduced the error between CFD and the GKN sensor. The orifice plate operates closer to its intended differential-pressure range, while the Pitot tube benefits from higher dynamic pressure.

Higher lift leads to better agreement too. Increasing lift enlarges the throttling area, raises mass flow, and improves both sensors' signal quality. For pressure ratio of 1.5 (with moderate absolute pressures) the error decreases almost linearly with lift as presented

in Figure 44. Passed a certain inlet pressure, the GKN sensor is already in its optimal range and lift has a weaker influence on accuracy as evident in Figure 46.

4.5.4 Temperature Prediction Errors

Temperature deviations grow with both pressure ration and lift, even while mass-flow errors shrink. This is likely due to no heat flux across the wall; with the adiabatic-wall boundary condition the CFD solver forbids any heat exchange between the gas and the surrounding metal. At higher mass-flow rates the flow velocity rises, and the core stream therefore remains colder than in reality, leading to a systematic under-prediction of bulk temperature. To quantify the influence of the wall model, one representative case was repeated with isothermal walls fixed to the inlet total temperature. Replacing the adiabatic condition with the isothermal one increased the predicted temperature from 271.33 K to 275.24 K, reducing the discrepancy and confirming that neglecting wall heat transfer is a source of error as shown in Figure 60. However, due to lack of experimental measurements of cylinder wall temperature, the iso-thermal wall boundary condition was not employed any further. A good step forward would be multi-region simulations involving conjugate heat transfer of solid material as explained further in the next section.




Figure 60. Choosing isothermal boundary condition shows improvement in thermal prediction. The simulated temperature has been raised by 3.91 K compared to adiabatic boundary condition.

4.5.5 Recommendations and Future Perspectives

1. **Conducting a comprehensive mesh sensitivity campaign:** a crucial next step is to perform a formal mesh sensitivity analysis to rigorously quantify the impact of discretisation error on the simulation outcomes. This thesis adopted a "proactive quality control" approach by using a single, acceptable quality mesh, which was a pragmatic choice to manage computational costs. However, this approach did not formally quantify the uncertainty introduced by the mesh itself. A formal grid convergence study, involving simulations on at least three systematically refined meshes, is a cornerstone of the verification and validation process. This will enable the explicit separation of discretisation error from modelling error (i.e., turbulence model inaccuracies), thereby establishing a clear "error budget." Ultimately, this analysis will provide critical guidelines on the optimal balance between computational cost and numerical accuracy, informing the structure of future DoE campaigns and increasing the overall confidence in the numerical predictions.
2. **Performing LES benchmarks:** to resolve the ambiguity in the sources of error, particularly in challenging flow regimes, it is recommended to perform high-fidelity LES benchmarks. The validation study revealed that the largest discrepancies between numerical and experimental results occurred at low-pressure ratios and low valve lifts, precisely where the two experimental sensors also showed the most disagreement. While this points towards experimental inaccuracies, the limitations of the RANS models in such complex flows cannot be dismissed. LES can serve as a powerful diagnostic tool to disentangle these convoluted error sources. By running an LES simulation, which is less dependent on modelling assumptions for large-scale turbulence, we can establish a high-fidelity numerical reference. If LES provides a better match to the experiment, the deviation between the RANS and LES results can be largely considered to be the RANS turbulence modelling error. The remaining deviation between LES and the experiment would then provide a much clearer estimate of the true experimental uncertainty,

allowing for a more targeted approach to improving both the physical test rig and the computational models.

3. **Refining the experimental setup and data acquisition:** the validation process revealed that experimental uncertainty was a dominant factor limiting the agreement between simulation and reality. Therefore, it is strongly recommended to refine the experimental setup for future campaigns. This involves two key actions: first, engaging with GKN to obtain the precise calibration envelope of their orifice meter to ensure comparisons are only made within its valid operating range. Second, and more critically, the inlet pipe section housing the Wärtsilä Pitot-static tube should be redesigned to meet established best-practice criteria (e.g., ≥ 7 pipe diameters of straight run). The current suboptimal placement downstream of a bend exposes the sensor to swirl and flow profile distortions, compromising data quality. Improving the fidelity of the experimental reference data is the single most effective way to enhance the entire validated workflow, as it provides a more reliable and unambiguous target for validating the numerical models.
4. **Extending the analysis to include Conjugate Heat Transfer (CHT):** The results from the pressure ratio sweep campaign provided clear evidence of a systematic deficiency in the thermal model: as the pressure ratio and mass flow increased, the temperature prediction error grew consistently. This trend strongly suggests that the adiabatic wall assumption is the primary cause. To test this hypothesis and improve thermal accuracy, it is recommended to extend the analysis to include CHT. By implementing CHT simulation for the high-pressure-ratio cases and calibrating it with cylinder-wall temperature measurements, we can explicitly quantify the improvement in thermal prediction. This will accurately capture the critical heat flux between the expanding gas and the surrounding metal components. Success in this area will not only reduce the temperature deviation but also refine the mass flow predictions through more accurate density calculations, ultimately providing a validated tool to assess the thermal loading on engine components, a vital consideration for the durability and performance of next-generation designs.

5. **Evaluating alternative mass flow measurement devices:** to definitively resolve the observed discrepancies between the two existing sensors, future campaigns should investigate the utility of an alternative mass flow measurement technology. Both the orifice plate and the Pitot tube are differential-pressure-based instruments, making them susceptible to similar sources of error related to fluid density and flow profile. Introducing a device like a Coriolis meter, which measures mass flow directly through inertial effects and is largely immune to these factors. This instrument could help reduce experimental uncertainties. This would, in turn, enable a far more rigorous assessment of the CFD model's predictive capabilities, allowing for a clearer distinction between numerical modelling errors and experimental artifacts.
6. **Extending validation to untested configurations and flow regimes:** the validated workflow should now be extended to include the cylinder head 2 and scavenging test points, which were beyond the scope of this thesis. Analysing the unique CH2 geometry and the complex cross-flow dynamics of scavenging is essential for a complete understanding of gas exchange. The robust PyFlowBench toolkit makes this extension efficient, paving the way for a comprehensive digital twin of the W31's breathing characteristics.
7. **Exploring Machine Learning applications for the numerical flow bench:** the successful automation and validation of the flow bench simulations via the PyFlowBench package is a critical enabler for the next frontier: the practical deployment of Machine Learning (ML). While the current automated workflow is efficient, the inherent computational cost of 3D CFD still limits the scale of design exploration. By leveraging the validated workflow to generate a rich training dataset, it is possible to develop models capable of predicting performance outcomes several orders of magnitude faster. This would make it feasible to iterate through hundreds, if not thousands, of design variations, exploring different cylinder head geometries and operating points in a way that is simply not possible with CFD alone. The significant speed-up, potentially reducing computational costs by multiple orders of magnitude (Caron et al., 2025; Hsieh & Felby, 2017),

directly supports the ultimate goal of optimising the gas exchange process in W31 engines, leading to better breathing, more complete combustion, enhanced fuel efficiency, and cleaner emissions. This thesis has laid the essential groundwork, proving the viability of the automated data pipeline required for such advanced, ML-driven development at Wärtsilä.

5 Conclusion

This thesis successfully achieved its primary objective: the development and validation of a robust, automated workflow for numerical flow bench simulations. The key to this achievement was the creation of the PyFlowBench toolkit, which automates the entire simulation pipeline from mesh generation to post-processing. Using this tool, an extensive validation campaign was conducted against experimental data from the Wärtsilä 31 flow bench. This process not only confirmed the predictive accuracy of the CFD model across various conditions but also demonstrated the workflow's power as a diagnostic tool, effectively identifying and quantifying key experimental uncertainties. The following sections summarise how this work addressed the guiding research questions.

1. How effective is workflow automation in enhancing the efficiency, productivity and consistency of engine-relevant CFD simulations?

PyFlowBench critically enabled the automation, consistency, and scalability of the entire CFD workflow. By reducing manual effort and ensuring repeatability from mesh generation to post-processing, it directly addressed RQ1, proving that structured automation significantly enhances the efficiency and reliability of engine-relevant CFD simulations. Crucially, PyFlowBench is also poised to be instrumental for future DoE and rapid prototyping endeavours, supporting advanced analysis.

2. To what extent can numerical simulations accurately replicate experimental data across various flow bench configurations and operating conditions?

Numerical simulations closely matched highly accurate experimental data, with the vast majority exhibiting deviations of approximately $\pm 5\%$. Yet, for lower pressure ratios, the numerical model's accuracy could be improved by including conjugate heat transfer, as the omission of wall temperature effects led to increased temperature deviations.

3. What factors primarily influence the uncertainty budget in non-reactive, transient valve flow CFD simulations and their corresponding flow-bench experiments? Also, how may these sources of uncertainty be identified, quantified, and possibly mitigated?

Thanks to the systematic CFD simulations, it became clear that there are considerable experimental inaccuracies. Thus, CFD can be used a driver to determine experimental uncertainties. On the experimental side, choosing the right measurement device suitable for the intended flow regime, calibration of such device and its location were identified as the source of accuracy or lack thereof. Numerically, accuracy is impacted by the appropriateness of numerical schemes for the flow type and mesh quality, the limitations of the chosen turbulence model, and the accurate representation of real-world conditions such as, wall temperature.

4. What potential do numerical simulations have for revealing underlying physical features that are challenging to capture through experimental measurements?

Numerical simulations allowed for visualisation and analysis of complete flow movement throughout the entire domain, consequently, tumble motion, swirls, recirculation zones, and high-velocity regions were identified. Experimentally capturing such detailed, volumetric flow information would require highly sophisticated setups, such as particle image velocimetry or similar advanced techniques, which are often limited in scope and resolution compared to a full CFD solution. By providing this deeper insight into the flow dynamics, numerical simulations help in improving future designs to minimise unwanted effects and optimise performance.

By addressing the aforementioned questions, this thesis successfully established and validated an automated CFD workflow for flow bench simulations, concurrently identifying

key experimental and numerical shortcomings. This foundational work lays the groundwork for significant future expansions, including mesh sensitivity analysis, large eddy simulation and conjugate heat transfer benchmarking, and exploring machine learning applications in flow bench simulations. It also enables further comprehensive studies on CH2 and CH3, along with scavenging test points, which remain relatively unexplored.

References

- A *Comprehensive Tour of snappyHexMesh*. (2012). <https://openfoamwiki.net/images/f/f0/Final-AndrewJacksonSlidesOFW7.pdf>
- Agnew, D. D. (1994, December 1). *What is Limiting Your Engine Air Flow: Using Normalized Steady Air Flow Bench Data*. <https://doi.org/10.4271/942477>
- Anderson, J. D. (2009). Governing Equations of Fluid Dynamics. In *Computational Fluid Dynamics* (pp. 15–51). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-85056-4_2
- Ashton, N., Unterlechner, P., & Blacha, T. (2018, April 3). *Assessing the Sensitivity of Hybrid RANS-LES Simulations to Mesh Resolution, Numerical Schemes and Turbulence Modelling within an Industrial CFD Process*. <https://doi.org/10.4271/2018-01-0709>
- Atkinson, C. (2014, September 30). *Fuel Efficiency Optimization using Rapid Transient Engine Calibration*. <https://doi.org/10.4271/2014-01-2359>
- Bhowmik, P. K., & Suh, K. Y. (2021). Flow mapping using 3D full-scale CFD simulation and hydrodynamic experiments of an ultra-supercritical turbine's combined valve for nuclear power plant. *International Journal of Energy and Environmental Engineering*, 12(3), 365–381. <https://doi.org/10.1007/s40095-021-00394-0>
- bp Energy Outlook 2024*. (2024).
- Bush, R. H., Chyczewski, T. S., Duraisamy, K., Einfeld, B., Rumsey, C. L., & Smith, B. R. (2019, January 7). Recommendations for Future Efforts in RANS Modeling and Simulation. *AIAA Scitech 2019 Forum*. <https://doi.org/10.2514/6.2019-0317>
- Caron, C., Lauret, P., & Bastide, A. (2025). Machine Learning to speed up Computational Fluid Dynamics engineering simulations for built environments: A review. *Building and Environment*, 267, 112229. <https://doi.org/10.1016/j.buildenv.2024.112229>
- Cerdoun, M., Khalfallah, S., Beniaiche, A., & Carcasci, C. (2020). Investigations on the heat transfer within intake and exhaust valves at various engine speeds. *International Journal of Heat and Mass Transfer*, 147, 119005. <https://doi.org/10.1016/j.ijheatmasstransfer.2019.119005>

- Chen, H., & Sick, V. (2017). Three-Dimensional Three-Component Air Flow Visualization in a Steady-State Engine Flow Bench Using a Plenoptic Camera. *SAE International Journal of Engines*, *10*(2), 2017-01–0614. <https://doi.org/10.4271/2017-01-0614>
- Choate, R., Schmaltz, K., Howard, S., Graves, P., Martin, J., & Kruckenberg, C. (2007). Air Flow Test Bench: A Senior Capstone Project. *2007 Annual Conference & Exposition Proceedings*, 12.192.1-12.192.12. <https://doi.org/10.18260/1-2--1544>
- Clenci, A. C., Iorga-Simăn, V., Deligant, M., Podevin, P., Descombes, G., & Niculescu, R. (2014). A CFD (computational fluid dynamics) study on the effects of operating an engine with low intake valve lift at idle corresponding speed. *Energy*, *71*, 202–217. <https://doi.org/10.1016/j.energy.2014.04.069>
- Curran, S., Onorati, A., Payri, R., Agarwal, A. K., Arcoumanis, C., Bae, C., Boulouchos, K., Dal Forno Chuahy, F., Gavaises, M., Hampson, G. J., Hasse, C., Kaul, B., Kong, S.-C., Kumar, D., Novella, R., Pesyridis, A., Reitz, R., Vaglieco, B. M., & Wermuth, N. (2024). The future of ship engines: Renewable fuels and enabling technologies for decarbonization. *International Journal of Engine Research*, *25*(1), 85–110. <https://doi.org/10.1177/14680874231187954>
- Cutrone, L., Schettino, A., Cardesa, J. I., Delattre, G., Coder, J. G., Qiang, S., Choudhari, M. M., & Vogel, E. (2024, January 8). Transition prediction in hypersonic regime on complex geometries with RANS-based models. *AIAA SCITECH 2024 Forum*. <https://doi.org/10.2514/6.2024-0291>
- Demirkesen, C., Colak, U., Savci, I. H., & Zeren, H. B. (2020). Experimental and Numerical Investigation of Air Flow Motion in Cylinder of Heavy Duty Diesel Engines. *Journal of Applied Fluid Mechanics*, *13*(2), 537–547. <https://doi.org/10.29252/jafm.13.02.30369>
- Domagala, M., & Fabis-Domagala, J. (2023). A Review of the CFD Method in the Modeling of Flow Forces. *Energies*, *16*(16), 6059. <https://doi.org/10.3390/en16166059>
- Eckl, R., Johnson, C., Shields, S., Cullipher, B., & Helmer, W. (2008). An Inexpensive Flow-Bench for Mechanical Engineering Labs. *ASEE Midwest Section Conference Proceedings*. <https://doi.org/10.18260/1-2-1153-51637>

- El-Adawy, M., Heikal, M. R., Aziz, A. R. A., Siddiqui, M. I., & Abdul Wahhab, H. A. (2017). Experimental study on an IC engine in-cylinder flow using different steady-state flow benches. *Alexandria Engineering Journal*, 56(4), 727–736. <https://doi.org/10.1016/j.aej.2017.08.015>
- Fang, L. (2025). *Introduction to Plenoptic Imaging* (pp. 1–6). https://doi.org/10.1007/978-981-97-6915-5_1
- Ferziger, J. H., Perić, M., & Street, R. L. (2020). *Computational Methods for Fluid Dynamics*. Springer International Publishing. <https://doi.org/10.1007/978-3-319-99693-6>
- Figliola, R. S., & Beasley, D. E. (2020). *THEORY AND DESIGN FOR MECHANICAL MEASUREMENTS*. JOHN WILEY.
- Flow and Level Measurement*. (n.d.). Retrieved June 27, 2025, from https://www.omega.co.uk/about/transactions/Transactions_Vol_IV.pdf
- Flow Performance*. (n.d.).
- Gowthama Krishnan, M., Rajkumar, S., Devarajan, Y., & Rajiv, A. (2024). A comprehensive review on advancement and challenges of renewable biofuelled reactivity controlled compression ignition (RCCI) engine. *Journal of the Energy Institute*, 113, 101540. <https://doi.org/10.1016/j.joei.2024.101540>
- Greenshields, C. (2024). *OpenFOAM v12 User Guide*. <https://doc.cfd.direct/open-foam/user-guide-v12>
- Hartmann, F., Buhl, S., Gleiss, F., Barth, P., Schild, M., Kaiser, S. A., & Hasse, C. (2016). Spatially Resolved Experimental and Numerical Investigation of the Flow through the Intake Port of an Internal Combustion Engine. *Oil & Gas Science and Technology – Revue d'IFP Energies Nouvelles*, 71(1), 2. <https://doi.org/10.2516/ogst/2015022>
- Hassan, M. A. S. M., Shahrman, A. B., Razlan, Z. M., Kamarrudin, N. S., Aziz, I. A., Khairunizam, W., Hashim, M. S. M., Harun, A., Ibrahim, I., Ibrahim, Z., Faizi, M. K., Rani, M. F. H., & Murali, R. (2021). *Engine performance enhancement by improving heat transfer in between exhaust valve and valve seat through CFD (transient thermal) simulation*. 020222. <https://doi.org/10.1063/5.0048421>

- Hsieh, C. C., & Felby, C. (2017). *AI Meets Physics - Altair physicsAI*. <https://www.ieabio-energy.com/wp-content/uploads/2018/02/Marine-biofuel-report-final-Oct-2017.pdf>
- ITABAR-Sensor Type IBR Operating Manual*. (n.d.).
- Jeong, S.-J., Kang, J., Moon, S.-J., & Lee, G. (2024). Transient and Dynamic Simulation of the Fluid Flow through Five-Way Electric Coolant Control Valve of a 100 kW Fuel Cell Vehicle by CFD with Moving Grid Technique. *Actuators*, 13(3), 110. <https://doi.org/10.3390/act13030110>
- Kim, D.-S., Yoo, Y.-J., Cho, Y.-S., & Ohm, I.-Y. (2003, May 19). *Numerical Analysis of Flow Characteristics for Intake Valve Design*. <https://doi.org/10.4271/2003-01-2015>
- Kiouranakos, K. I., de Vos, P., Zoumpourlos, K., Coraddu, A., & Geertsma, R. (2025). Methanol for heavy-duty internal combustion engines: Review of experimental studies and combustion strategies. *Renewable and Sustainable Energy Reviews*, 214, 115529. <https://doi.org/10.1016/j.rser.2025.115529>
- Kudźma, Z., & Stosiak, M. (2015). Studies of flow and cavitation in hydraulic lift valve. *Archives of Civil and Mechanical Engineering*, 15(4), 951–961. <https://doi.org/10.1016/j.acme.2015.05.003>
- Ma, F., Guo, L., Li, Z., Zeng, X., Zheng, Z., Li, W., Zhao, F., & Yu, W. (2023). A Review of Current Advances in Ammonia Combustion from the Fundamentals to Applications in Internal Combustion Engines. *Energies*, 16(17), 6304. <https://doi.org/10.3390/en16176304>
- Menter, F. R. (1994). Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal*, 32(8), 1598–1605. <https://doi.org/10.2514/3.12149>
- Moffat, R. J. (1988). Describing the uncertainties in experimental results. *Experimental Thermal and Fluid Science*, 1(1), 3–17. [https://doi.org/10.1016/0894-1777\(88\)90043-X](https://doi.org/10.1016/0894-1777(88)90043-X)
- Mohamad Shafie, N. A., Muhamad Said, M. F., Abdul Aziz, A., Abdul Latiff, Z., Mat Yamin, A. K., & Tamaldin, N. (2017). Steady flow and dynamic analyses comparison of an air intake breathing capacity. *MATEC Web of Conferences*, 90, 01067. <https://doi.org/10.1051/mateconf/20179001067>

- Mohammadebrahim, A., Shafiei, B., & Kazemzadeh Hannani, S. (2012). Numerical simulation of in-cylinder tumble flow field measurements and comparison to experimental results. *The Journal of Engine Research*.
- Montorfano, A. (2017). *Mesh generation for HPC problems: the potential of SnappyHex-Mesh*.
- Moujaes, S. F., & Jagan, R. (2008). 3D CFD Predictions and Experimental Comparisons of Pressure Drop in a Ball Valve at Different Partial Openings in Turbulent Flow. *Journal of Energy Engineering*, 134(1), 24–28. [https://doi.org/10.1061/\(ASCE\)0733-9402\(2008\)134:1\(24\)](https://doi.org/10.1061/(ASCE)0733-9402(2008)134:1(24))
- Nishad, K., Ries, F., Li, Y., & Sadiki, A. (2019). Numerical Investigation of Flow through a Valve during Charge Intake in a DISI -Engine Using Large Eddy Simulation. *Energies*, 12(13), 2620. <https://doi.org/10.3390/en12132620>
- Noh, H. K., & No, S.-Y. (2017). Effect of bioethanol on combustion and emissions in advanced CI engines: HCCI, PPC and GCI mode – A review. *Applied Energy*, 208, 782–802. <https://doi.org/10.1016/j.apenergy.2017.09.071>
- Reitz, R. D., Ogawa, H., Payri, R., Fansler, T., Kokjohn, S., Moriyoshi, Y., Agarwal, A., Arcoumanis, D., Assanis, D., Bae, C., Boulouchos, K., Canakci, M., Curran, S., Denbratt, I., Gavaises, M., Guenther, M., Hasse, C., Huang, Z., Ishiyama, T., ... Zhao, H. (2020). IJER editorial: The future of the internal combustion engine. *International Journal of Engine Research*, 21(1), 3–10. <https://doi.org/10.1177/1468087419877990>
- Roelofs, F., & Shams, A. (2019). CFD—Introduction. In *Thermal Hydraulics Aspects of Liquid Metal Cooled Nuclear Reactors* (pp. 213–218). Elsevier. <https://doi.org/10.1016/B978-0-08-101980-1.00006-5>
- Singh, M., Singla, M. K., Beryozkina, S., Gupta, J., & Safaraliev, M. (2024). Hydrogen vehicles and hydrogen as a fuel for vehicles: A-State-of-the-Art review. *International Journal of Hydrogen Energy*, 64, 1001–1010. <https://doi.org/10.1016/j.ijhydene.2024.03.325>
- Smulter, B. (2024). *W31 Cylinder head flow measurements at GKN*.

- Soner, S., & Ozturan, C. (2015). Generating Multibillion Element Unstructured Meshes on Distributed Memory Parallel Machines. *Scientific Programming*, 2015, 1–10. <https://doi.org/10.1155/2015/437480>
- Soriano, B. S., & Rech, C. (2012). *STEADY DISCHARGE COEFFICIENT IN INTERNAL COMBUSTION ENGINE*. <https://abcm.org.br/anais/encit/2012/links/pdf/ENCIT2012-158.pdf#:~:text=The%20intake%20system%20is%20efficient,Ferguson%2C%201985>
- Taghavi, R., Dupont, A., & Dupont, J. F. (1990). Aerodynamic and Thermal Analysis of an Engine Cylinder Head Using Numerical Flow Simulation. *Journal of Engineering for Gas Turbines and Power*, 112(3), 335–340. <https://doi.org/10.1115/1.2906500>
- Theodorakakos, A. (2024). Numerical Simulation and Comparison of Different Steady-State Tumble Measuring Configurations for Internal Combustion Engines. *Computation*, 12(7), 138. <https://doi.org/10.3390/computation12070138>
- Valdés, J. R., Rodríguez, J. M., Monge, R., Peña, J. C., & Pütz, T. (2014). Numerical simulation and experimental validation of the cavitating flow through a ball check valve. *Energy Conversion and Management*, 78, 776–786. <https://doi.org/10.1016/j.enconman.2013.11.038>
- Wahono, B., Jwa, K., & Lim, O. (2019). A Study on In-Cylinder Flow of Small Engine Using Steady-State Flow Benches. *Energy Procedia*, 158, 1856–1862. <https://doi.org/https://doi.org/10.1016/j.egypro.2019.01.432>
- Wang, Z., Shuai, S., Li, Z., & Yu, W. (2021). A review of energy loss reduction technologies for internal combustion engines to improve brake thermal efficiency. In *Energies* (Vol. 14, Issue 20). MDPI. <https://doi.org/10.3390/en14206656>
- Welch, C., Illmann, L., Schmidt, M., & Böhm, B. (2023). Experimental characterization of the turbulent intake jet in an engine flow bench. *Experiments in Fluids*, 64(5), 91. <https://doi.org/10.1007/s00348-023-03640-9>
- White, F. M. (2011). *Fluid Mechanics*.
- Yang, X., Kuo, T.-W., Guralp, O., Grover, R. O., & Najt, P. (2017). In-Cylinder Flow Correlations Between Steady Flow Bench and Motored Engine Using Computational Fluid

Dynamics. *Journal of Engineering for Gas Turbines and Power*, 139(7).
<https://doi.org/10.1115/1.4035627>

Appendices

Appendix 1. Representative Code Snippets from PyFlowBench Scripts

A snippet of SnappyHexMesh script:

```
# Choose your queue and number of cores:
#$ -pe orte 108
...
# Set flow configuration
Configuration="$1"
IV1="$2"
IV2="$3"
EV1="$4"
EV2="$5"
...

# Move the intake and exhaust valves for the given IV1, IV2, EV1,
EV2 lift values
for f in constant/geometry/*_IV1*.stl; do surfaceTransformPoints
"translate=(0 0 $IV1)" "$f" "${f%.stl}.stl"; done
for f in constant/geometry/*_IV2*.stl; do surfaceTransformPoints
"translate=(0 0 $IV2)" "$f" "${f%.stl}.stl"; done
for f in constant/geometry/*_EV1*.stl; do surfaceTransformPoints
"translate=(0 0 $EV1)" "$f" "${f%.stl}.stl"; done
for f in constant/geometry/*_EV2*.stl; do surfaceTransformPoints
"translate=(0 0 $EV2)" "$f" "${f%.stl}.stl"; done
```

Code 1. A Bash script snippet from the PyFlowBench workflow, demonstrating the automated translation of valve STL surfaces based on command-line lift arguments.

In this snippet, the script first specifies the computational resources by defining the number of HPC cores allocated for the meshing job. It then assigns user-supplied arguments to variables that define the flow configuration (intakeFlow, exhaustFlow, or scavenging) and the corresponding intake and exhaust valve lifts. Using these parameters, the script dynamically adjusts the STL geometry files by translating the valve components along the Z-axis to match the specified lift values. This automated transformation of geometry ensures that each simulation case is accurately parameterised, enabling rapid and consistent mesh generation across multiple cases, a critical enabler for high-throughput CFD studies.

A snippet of data transfer automation script:

```

def _find_column(df: pd.DataFrame, patterns: Sequence[str],
                 *, label: str) -> str:
    normals = {_normalise(c): c for c in df.columns}
    for pattern in patterns:
        if pattern in normals:
            return normals[pattern])

...

COLS_COMMON = {
    "temperature": ["inletairtemp/ave"],
    "wartsila": ["massflow/ave"]
    "gkn": ["massflow_gkn/ave"],
    "pitot": ["inletairpitot/ave"],
    "dp_nozzle": ["dp_pr_nozzle_gkn"],
    "ch_temp": ["ch temp"]}

COLS_INTAKE = {
    "valve": ["iv\nvalvelift"],
    "p_inlet": ["ca pipe end ave abs"],
    "p_outlet": ["ch pr downstream abs"]}

COLS_EXHAUST = {
    "valve": ["ev\nvalvelift", "evlift"],
    "p_inlet": [ "ch pr downstream abs"],
    "p_outlet": "exh pipe end pr1 abs"]}

...

df_full = pd.read_excel(
    DATASET_FILE,
    sheet_name=sheet_name,
    engine="openpyxl",
    header=header_row, )

...

def _update_dictionary(dictionary_text: str, *, scenario:
                       str, df: pd.DataFrame) -> str:
    is_intake = scenario == "in"
    cols = {**COLS_COMMON, **(COLS_INTAKE if is_intake else
                               COLS_EXHAUST)}

...

# ----- Meta -----
dictionary_text = (
    dictionary_text.replace(
        "PLACEHOLDER_SIMDIR",
        "sim_intk_validation_data" if is_intake else

```

```

        "sim_exhst_validation_data",
    )
    .replace("PLACEHOLDER_CASETYPE", "intakeFlow" if
            is_intake else "exhaustFlow")

# ----- Valve lifts -----
...
if is_intake:
    dictionary_text = replace_first(dictionary_text,
                                    "PLACEHOLDER_IV", lift_vals)
...

# ----- Pressures & temperature -----
dictionary_text = (
    dictionary_text.replace("PLACEHOLDER_PINLET", pinlet)
    .replace("PLACEHOLDER_POUTLET", poutlet)
    .replace("PLACEHOLDER_TEMPERATURE", temps))

# ---- Mass-flow rates (validity) -----

# ---- Cylinder-head temperature -----
...

return dictionary_text

```

Code 2. Python script for populating the flowBenchDict with experimental data.

In the above snippet, the experimental dataset is first read from an Excel file. Relevant columns are then identified using pattern matching, starting with those common to all configurations (e.g., temperature, mass flow rates), followed by configuration-specific columns depending on whether the scenario is intake or exhaust. These extracted values are used to dynamically update the Flow benchDict by replacing predefined placeholders with the corresponding data, such as valve lifts, inlet/outlet pressures, and temperatures. This process ensures accurate and automated transfer of experimental data into the CFD configuration file, enabling consistent and error-free simulation setup.

Pre-processing automation and simulation submission:

```

def _checks_for_boundary_conditions(self):
    ...
    if self.case_type == "intakeFlow":
        if not(
            len(self.valve_lift_iv [1:]) == len(self.p_inlet) == len(self.p_outlet) ==
            len(self.temperature) ==

```

```

        len(self.wartsilaMassFlowRate)                ==
        len(self.gknMassFlowRate)
    ):
        raise ValueError(
            'The number of entries for boundary conditions,
            Mass Flow Rate and iv lifts should be equal'
        )
...

def build_cases(self):
    """Build cases based on valve lift values."""
    self._checks_for_boundary_conditions()

    ...

    self.case_name = f"{self.case_type}_{str(lift)}"
    print(f"Creating case: {self.case_name}")
    case_dir = self.create_case_directory(self.case_name)
    self.copy_lift_mesh_snappy(case_dir, self.case_type,
        lift, dryrun=False)
    # Conditionally execute merge_mesh_modules based
    on snappy value

    if self.snappy == False:
        self.merge_mesh_modules(case_dir, lift)
        else:
            print("Skipping merge_mesh_modules due to snappy
            setting.")
            print("Running the solver 'foamRun' with a job
            script: ")
            try:
                os.chdir(self.case_dir)
                print(f"Successfully changed the working di
                    irectory to: {os.getcwd()}")
                job_name = job_prefix + str(lift)
                self.execute_commands('runCase', job_name)

def create_case_directory(self, case_name):
    """Creates a case directory using foamCloneCase in
    the run directory."""
    # Create the case directory for this simulation case
    self.case_dir = os.path.join(self.sim_dir, case_name)
    flow_dict = 'constant/Flow benchDict'
    print(f"case directory: {self.case_dir}")
    if not os.path.exists(self.case_dir):
        try:
            # Clone base case
            subprocess.run(['foamCloneCase',
                self.base_dir, self.case_dir], check=True)
            print(f"Successfully created case directory:
                {self.case_dir}")
            ...
        return self.case_dir

```

```

...
def execute_commands(self, script_name, job_name, dry_run=False):
    try:
        # Submit the script to the job scheduler
        qsub_command = ['qsub', '-N', job_name, script_name]
        if dry_run:
            ...
        else:
            result = subprocess.run(qsub_command, check=True,
                                    capture_output=True, text=True)
            print(f"Submitted {script_name} to the job scheduler.")
            print(f"qsub output:\n{result.stdout}")

            # Monitor the job
            qstat_command = ['qstat']

            ...

```

Code 3. Core functions from the PyFlowBench package demonstrating the automated workflow for case validation, creation, and submission to an HPC job scheduler.

This script automates the preparation and submission of multiple simulation cases based on varying valve lifts and corresponding boundary conditions. It begins by validating that all necessary parameters, such as inlet/outlet pressures, temperatures, and mass flow rates, are consistently provided for each case using `_checks_for_boundary_conditions`. This prevents misaligned or incomplete configurations that could compromise simulation accuracy.

The `build_cases` function orchestrates the entire case setup process. It constructs uniquely named case directories using `foamCloneCase`, copies the appropriate mesh based on the valve lift, and handles either modular mesh merging or skips it based on the mesh type (e.g., `snappyHexMesh` vs. `GridPro`). It then calls the job submission routine.

Finally, `execute_commands` submits the simulation job to the HPC scheduler via `qsub`, optionally monitoring the job and printing real-time feedback. This automation significantly reduces manual effort, ensures consistent configuration across cases, and enables

scalable execution of high-volume simulation campaigns, which is fundamental for achieving both efficiency and reproducibility in CFD workflows.

Postprocessing scripts:

```
'''Mass flow rate calculation and plotting'''

# Path to Flow benchDict
flow_bench_dict_path = '/Flow benchDict'

# Read the Flow benchDict
data_dict = read_dict(Path(flow_bench_dict_path), py-
thon_types=True)

# Extract simulation directory
sim_dir = data_dict.get('simDir')

# Extract Boundary conditions
turbulence_model = ...
case_type
inlet_pressure
outlet_pressure
wartsila_mass_flow_list
gkn_mass_flow_list

# Determining the mass flow rate location
if case_type == "exhaustFlow":
    file_path = os.path.join(case_dir, 'postProcessing',
    'patchFlowRate(patch=inlet)/0/surface
    FieldValue.dat')
    ...

# Load the data from the file
data = np.loadtxt(file_path, skiprows=4)

# Extract time and mass flow rate
time = data[:, 0]
mass_flow = np.abs(data[:, 1])

# Filter data after stabilization
valid_time = time >= 0.3
mass_flow_filtered = mass_flow[valid_time]

# Compute time-averaged mass flow rate
avg_mass_flow = np.mean(mass_flow_filtered)

# Calculate error percentages (if experimental values
                                exist)
percent_error = ((avg_mass_flow - exp_flow_wartsila) /
```

```

exp_flow_wartsila * 100) if exp_flow_wartsila else None

error_gkn = ((avg_mass_flow - exp_flow_gkn) / exp_flow_gkn *
100) if exp_flow_gkn else None

# Plotting
plt.figure(figsize=(10, 6))
...

'''Temperature calculation and plotting'''
# Define the path to the temperature probe data file
temp_path = os.path.join(case_dir, 'postProcessing',
                        'probe_temperature', '0', 'T')

# Load the temperature probe data (skip header lines)
try:
    temp_data = np.loadtxt(temp_path, skiprows=2)

# Extract time and temperature values
    time = temp_data[:, 0]
    temp = temp_data[:, 1]

# Filter the data for time values after stabilization
valid = time >= 0.2
avg_temp = np.mean(temp[valid])

# Retrieve the experimental temperature for this lift (now in
Kelvin)
exp_temp = exp_temperature_list[idx] if idx < len(exp_tempera-
ture_list) else None
temp_error = ((avg_temp - exp_temp) / exp_temp * 100) if exp_temp
is not None else None

# Plot the temperature data over time with comparison to the
experimental value
plt.figure(figsize=(10, 5))

```

Code 4. Python script from the PyFlowBench package for the automated post-processing and validation of simulation results.

This postprocessing script automates the extraction, analysis, and visualisation of critical CFD output quantities, mass flow rate and temperature, directly from the simulation result files. For mass flow rate, the script identifies the correct output file based on the flow configuration (e.g., exhaustFlow) and reads time-resolved flow rate data. It filters out the transient startup phase and calculates a time-averaged mass flow value after stabilization. This value is then compared against experimental benchmarks (e.g., from

Wärtsilä or GKN sensors), and percentage errors are computed. The results are also plotted for visual assessment and reporting.

For temperature, a similar procedure is followed: probe data is read, filtered post-stabilization, averaged, and compared against measured cylinder head temperatures. Visualisation tools such as Matplotlib are used to plot both simulated and experimental trends for validation.

Cleanup and memory management script:

```
read -rp $'Enter the path to the simulation directory (e.g.
sim_intk_validation_valve_sweep):\n> ' CASE_DIR

cd "$CASE_DIR"

PS3=$'Choose reconstruction mode for all cases (enter number):'
select CHOICE in \
  "Run \`reconstructPar\`" \
  "Run \`reconstructPar -latestTime\`" \
  "Skip all cases"; do
  case $REPLY in
    1) CMD=(reconstructPar);          break ;;
    2) CMD=(reconstructPar -latestTime);  break ;;
    3) CMD=();                        break ;;
    *) printf 'Invalid selection. Try again.\n' ;;
  esac
done

...

if ((${#CMD[@]})); then
  printf 'Running: %s\n' "${CMD[*]}"
  if "${CMD[@]}"; then
    printf 'Removing processor directories...\n'
    rm -rf processor*
  fi
fi

...

if [[ -d standardized_meshes ]]; then
  printf '\nRemoving standardized_meshes directory...\n'
  rm -rf standardized_meshes
fi

printf '\nAll cases processed. Done.\n'
```

Code 5. A utility Bash script for the post-simulation reconstruction and cleanup of parallel OpenFOAM cases.

This Bash script provides an interactive, automated routine for final cleanup and memory management after CFD simulations have been executed. Upon execution, the user is prompted to enter the path to the simulation directory. The script then allows the user to choose between two modes of field data reconstruction:

Full reconstruction of all timesteps (`reconstructPar`)

Reconstruction of only the latest timestep (`reconstructPar -latestTime`)

Or skipping reconstruction entirely

Once reconstruction is completed (if selected), the script automatically deletes the `processor*` directories that were created for parallel computation. These directories can be large in size and are no longer needed post-reconstruction. It also checks for and removes the `standardized_meshes` directory, if present, to further reduce unnecessary storage.

This script plays a crucial role in memory optimisation and disk space management, especially when handling large-scale simulation campaigns across multiple cases. By embedding interactive controls and safety checks, it ensures controlled cleanup while freeing up high-volume HPC storage.

Appendix 2. Uncertainty Analysis for Mass Flow Rate Measurement

This appendix provides a detailed, step-by-step calculation of the propagated uncertainty for the experimental mass flow rate, as summarised in the main text. The analysis follows the single-sample uncertainty methodology of Moffat (1988) using the RSS combination of individual error sources. Sensitivity coefficients are determined numerically using the sequential perturbation method. All uncertainties are specified at a 95% confidence interval.

In order to assess the uncertainties in the measured mass flow rate (\dot{m}), the governing physical equations need to be considered for the primary measurements constituting the final mass flow rate. The ITABAR sensor is an averaging Pitot tube, for which the manufacturer's documentation provides the governing equation for mass flow (\dot{m}) as follows:

$$\dot{m} = k * 1.0159 * A * \sqrt{2\rho\Delta P}. \quad (29)$$

In Equation 29, K is the empirical K-Factor of the sensor, A is the cross-sectional area of the pipe, ρ is the fluid density at operating conditions, ΔP is the differential pressure measured by the sensor, and 1.0159 correction factor, which accounts for compressibility and other calibration constants specific to the ITABAR system. Also using the Ideal Gas Law presented in Equation (16), the density of the working fluid (air) is measured.

Combining these principles, it is clear that the final calculated mass flow rate, \dot{m} , is a function of several independent, measured variables: $\dot{m} = f(K, A, \Delta P, P_1, T_1)$. The uncertainty in \dot{m} is therefore fundamentally dependent on the uncertainties in each of these inputs.

Baseline Operating Conditions

The baseline calculation uses the "Normal" operating point values from the Intra-Automation Winflow 3.1 sizing sheet as summarised in Table 12.

Table 12. A summary of baseline operating conditions.

Parameter	Symbol	Value
Pipe diameter	d	159.3 mm (0.1593 m)
Cross-sectional area	A	0.01993 m ²
Operating pressure	P_1	951,330 Pa
Operating temperature	T_1	291.15 K
Fluid density	ρ	11.376 kg/m ³
Differential pressure	ΔP	19,317 Pa
K-Factor	K	0.5967

The verified nominal mass flow rate is 8.00 kg/s.

Individual Variables and Estimated Uncertainties

The following Table 13 summarises the nominal values and estimated standard uncertainties ($k=1$, or 68% confidence) for each independent variable. As per standard practice, these standard uncertainties are then expanded to a 95% confidence interval ($k=2$) for the final result.

Table 13. Summary of estimated uncertainties.

Variable	Symbol	Value	Standard Uncertainty ($k=1$)	95% Uncertainty	Source
K-Factor	K	0.5967	± 0.00203 (0.34%)	± 0.00406 (0.68%)	Bias Error. Source: ITABAR Technical Catalogue
Pipe diameter	d	159.3 mm	± 0.5 mm	± 1.0 mm	Bias Error. Source: ITABAR Technical Catalogue
Differential Pressure	ΔP	19,317 Pa	± 13.65 Pa	± 27.3 Pa	Combined Error. From Aplsens datasheets.

Absolute Pressure	P_1	951,330 Pa	± 1250 Pa	± 2500 Pa	Assumed Error. Based on a standard industrial transmitter.
Temperature	T_1	291.15 K	± 0.093 K	± 0.186 K	Assumed Error. Based on a Class A Pt100 RTD

Sequential Perturbation and Contribution Analysis

Each variable is perturbed by its uncertainty interval (δX_i) to calculate its individual contribution (UX_i) to the total uncertainty in mass flow rate. While calculations could proceed using standard uncertainties ($k=1$) and be expanded at the final stage, the sequential perturbation method using the 95% interval is applied here for clarity, as the non-linear relationship of terms like differential pressure (under the square root) means the final result is largely insensitive to this choice. All calculations were performed using full floating-point precision to minimise rounding errors.

$$K_{new} = K_0 + \delta K = 0.5967 + 0.00406 = 0.60076$$

$$\dot{m}_{new} = 1.0159 * K_{new} * A_\varepsilon * \sqrt{2 * \rho_0 * \Delta P_0} = 8.0638$$

$$U_k = |\dot{m}_{new} - \dot{m}_0|$$

$$U_k = |8.06 - 8.00| = 0.0638 \frac{kg}{s}$$

The final calculated perturbations for the rest of the variables is presented below in Table 14.

Table 14. Perturbation results.

Variable	Perturbed Value	Resulting \dot{m} (kg/s)	Contribution (UX_i)
K	0.60076	8.0638	0.0638
d	0.1603 m	8.1105	0.1105
ΔP	19344.3 Pa	8.0149	0.0149
P_1	953,830 Pa	8.0222	0.0222
T_1	291.336 K	8.0091	0.0091

The total uncertainty is the root-sum-square of the individual contributions is as follows:

$$\delta_{\dot{m}} = \sqrt{U_k^2 + U_d^2 + U_{\Delta P}^2 + U_{P_1}^2 + U_{T_1}^2},$$

$$\delta_{\dot{m}} = \frac{0.13068416124kg}{s}.$$

The final result for the mass flow rate, expressed with its 95% confidence uncertainty, is: $\dot{m} = 8.00 \pm 0.131 \text{ kg/s}$. And the relative uncertainty is calculated as: $Relative\ uncertainty = \left(\frac{0.13}{8.00}\right) * 100 \% = \pm 1.6 \%$. In Table 15, the most important takeaway from this analysis has been summarised mentioning the final uncertainty contribution from each of the primary components.

Table 15. Final sensitivity analysis table.

Source Variable (X_i)	Uncertainty Contribution (U_{X_i})	Variance ($U_{X_i}^2$)	Contribution to Total Variance ($(U_{X_i}^2 / \delta_{\dot{m}}^2) * 100\%$)
Pipe diameter	0.11	0.0121	72.89 %
K-Factor	0.0638	0.00407	24.52 %
Differential Pressure	0.0149	0.0002	1.2 %
Absolute Pressure	0.0222	0.0004	2.4 %
Temperature	0.0091	0.00008	0.48%

This sensitivity analysis confirms that the total uncertainty is overwhelmingly dominated by two factors: the manufacturing tolerance of the pipe diameter (72.89%) and the calibration uncertainty of the ITABAR K-Factor (24.52%). Combined, these two sources account for almost 98% of the total uncertainty in the final mass flow rate measurement.