UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

UNIVERSITÀ DEGLI STUDI
DI MODENA E REGGIO EMILIA

Dottorato di ricerca in Ingegneria Industriale e del Territorio

Ciclo XXXVIII

CFD modelling of innovative spark ignition engines running on biofuels and hydrogen

Candidato: Claudiu Marcu Di Gaetano Iftene

Relatore (Tutor): Prof. Enrico Mattarelli

Coordinatore del Corso di Dottorato: Prof. Alberto Muscio

Acknowledgements

To my parents, who from my earliest years gave me the chance to study and become who I am.

To my tutors, who supported me throughout these activities.

To my “affettati”, who gave me peace of mind.

To my friends, who gave me a smile.

To myself, for never giving up.

Abstract

Reducing greenhouse-gas emissions from road transport requires practical pathways that improve efficiency and lower pollutant formation across diverse duty cycles. While electrification is expanding, internal-combustion engines (ICEs) remain critical in many applications due to energy density, operating range, and refueling logistics. In this setting, low-carbon fuels (e.g., ethanol blends, e-fuels, hydrogen) and unconventional architectures (e.g., opposed-piston) are attractive options, provided that their design and calibration are guided by predictive, economical simulation tools. This dissertation develops a reproducible workflow for three-dimensional computational fluid dynamics (3D CFD) that couples disciplined pre-processing and model selection with minimal, targeted calibration against experiments, aiming to deliver credible predictions at manageable cost.

Four studies structure the work. First, E85 is assessed against a gasoline baseline in a single-cylinder, four-stroke PFI engine using a premixed approach validated by 1D/3D correlation. Equivalence-ratio and spark-timing sweeps show that E85 favors leaner operation ($\varphi \approx 0.90–0.95$) and slightly earlier phasing, increasing combustion efficiency with only modest performance penalties after retuning; the reduced MFB10–90 is consistent with ethanol's higher laminar flame speed.

Second, a four-cylinder PFI engine with detailed chemistry confirms and refines these trends: two calibration levers—an energy-addition factor (EF) and a turbulence–chemistry interaction (TCI) multiplier—suffice to match experiments. Under the studied conditions, E85 yields higher IMEP, lower CO/HC and higher CO_2, and a reduced MFB10–90 relative to E5, at the cost of higher injected mass dictated by stoichiometry.

Third, a hydrogen-fueled opposed-piston spark-ignition engine is investigated in OpenFOAM using structured cylinder motion and custom post-processing utilities (volume/flux accounting; fresh/burnt species tagging). Cold-flow metrics reproduce limiting-model behavior and align with 1D trends. Combustion simulations at very lean mixtures ($\lambda \approx 2.4$) identify the expected phasing trade-off: IMEP peaks near 30° bTDC, while 35° bTDC improves efficiency without additional work.

Fourth, a DISI spray-model calibration campaign in STAR-CCM+ distills practical choices (source placement away from walls; Blob vs. Rosin–Rammler primary-breakup consistency; KH–RT tuning). Importantly, across the tested operating conditions, the calibrated spray setup reproduces experimental trends without modifying the previously defined calibration coefficients, demonstrating predictive robustness and transferability.

Overall, the thesis contributes a disciplined 3D CFD practice—spanning combustion and sprays, commercial and open-source tools—that shortens the path from concept to credible prediction and is readily extensible to other fuels, architectures, and operating regimes.

Contents

List of Figures

Introduction

Road transport underpins modern economies and daily life. As of 2020, more than one billion light-duty vehicles were in operation worldwide, and the fleet is expected to grow further [1]. While this scale has brought substantial socio-economic benefits, it has also raised concerns about energy use and greenhouse-gas (GHG) emissions: the transport sector accounts for roughly 15% of global energy-related GHG emissions [2].

Internal-combustion engines (ICEs) still power the vast majority of on-road and maritime fleets and remain a mature yet improvable technology. Despite rapid progress in electrification, ICE development remains relevant where batteries face constraints in cost, infrastructure, duty cycle, or energy density (e.g., long-haul, off-road, maritime, aviation) [3].

Current research to improve ICE performance and sustainability advances along complementary levers:

- Thermodynamics and combustion: mixture formation, ignition systems, and advanced modes such as lean/pre-chamber, low-temperature combustion, and dual-fuel strategies
- Architecture and hardware: boosting and air handling, variable compression ratio, friction reduction, unconventional layouts
- Fuels and aftertreatment: low-carbon fuels combined with catalytic systems
- System-level control and integration: hybridization, energy management, and calibration

Among low-carbon fuels, biofuels are produced from biomass feedstocks (e.g., agricultural residues, forestry by-products, algae) and are often classified by “generation” according to feedstock origin. Their sustainability depends on feedstock type, conversion efficiency, and land-use impacts, commonly framed as land-use change (LUC) and indirect LUC (ILUC) [3]. E-fuels (power-to-X, PtX) are synthetic fuels produced by combining hydrogen—typically made from renewable-electricity electrolysis—with captured CO_2, yielding molecules from methane and methanol to drop-in hydrocarbons (power-to-gas/liquid, PtG/PtL) [4]. In principle, efuels can approach a closed carbon loop on a well-to-wheel (WtW) basis when supplied by low-carbon electricity and sustainable CO_2 sources; however, they entail multi-step supply chains and significant infrastructure needs that reduce WtW efficiency. [5]. Quantitative comparisons among battery electric vehicles (BEVs), biofuels, and efuels must therefore be conducted on a transparent WtW basis with clearly stated assumptions (electricity mix, capture route, synthesis efficiency, duty cycle).

In parallel, unconventional engine architectures—most notably opposed-piston and free-piston configurations, often operating on a two-stroke cycle—are being investigated for their potential to reduce losses, enable uniflow scavenging, and broaden feasible combustion strategies (e.g., homogeneous-charge compression ignition, low-temperature combustion, reactivity-controlled compression ignition).

At the vehicle-system level, hybrid powertrains (series, parallel, or blended) couple ICEs with electric machines to improve real-world efficiency via load leveling and energy recovery.

These benefits come with added complexity in energy management, packaging, thermal control, and cost—factors that must be evaluated alongside environmental performance.

Non-exhaust emissions—e.g., tire and brake wear, lubricant consumption—also affect air quality. Recent regulatory frameworks, such as Euro 7, increasingly address these sources, underscoring the need for holistic vehicle design and assessment [6].

In summary, decarbonizing road transport requires a portfolio of solutions rather than a single pathway. This motivates a rigorous, WtW-consistent assessment of advanced combustion concepts, unconventional engine architectures, and low-carbon fuels, together with their integration into hybridized powertrains.

During my PhD, I worked on multiple projects involving ICEs. Specifically, I investigated biofuels (E25, E85), efuels and hydrogen (H_2) for ICE applications—covering gas exchange, combustion processes, flow characterization, and spray-bomb calibrations. These studies were conducted mainly with three-dimensional Computational Fluid Dynamics (3D CFD) tools, and I gained extensive experience with both commercial and open-source software (OpenFOAM 8).

Commercial 3D CFD tools provide robust and reproducible workflows. However, their modules are typically tailored to specific applications (e.g., conventional four-stroke ICEs) and are not readily applicable beyond those intended use cases. By contrast, open-source 3D CFD frameworks, with accessible source code, enable users to extend capabilities and model unconventional architectures, including opposed-piston and free-piston engines.

Here I present the most relevant results from my projects; taken together, they also trace the workflow through which I developed strong working knowledge of 3-D CFD. A common feature of these projects was the definition of a methodology to obtain the best result with the least possible resources—particularly in OpenFOAM 8—by identifying a specific routine for meshing, pre-processing, and simulation strategy.

This thesis presents three university projects conducted with 3-D CFD tools:

- E85 usage in single-cylinder, 4-stroke, PFI motorcycle ICE.
  This work had the focus to investigate the substitution of traditional gasoline with E85. Starting from experimentally calibrated gasoline model, E85 fuel was set up which results are compared with those from gasoline case. ECFM-3Z model was used to predict combustion behavior
- Gasoline and E85 combustion comparison in a 4-cylinder, 4-stroke, PFI ICE.
  Experimental data for the above mentioned engine were available for both fuels. 3D CFD model for both gasoline and E85 were built and calibrated. Combustion was modelled using detailed chemistry kinetics
- Virtual development of Opposed-piston engine fueled by hydrogen using OpenFOAM 8. In this project, methodologies for meshing, pre-processing, simulation and post-processing were developed. Combustion model was provided by Politecnico di Milano

During my research period abroad, I focused on calibrating spray models for Direct Injection Spark Ignition (DISI) applications. As sensitive information cannot be shared, this section offers a concise, non-proprietary summary of spray-correlation best practices.

Ethanol is a well-known alternative to conventional fossil fuels for ICEs. It can be used neat or blended with gasoline. For gasoline blends, anhydrous ethanol is required, typically <1% (v/v) water. Neat fuel ethanol can be hydrous (≈5% v/v water), as used in some markets (e.g., Brazil). Dehydration beyond the ethanol–water azeotrope requires additional processing and energy [7].

Ethanol's high latent heat of evaporation and high octane number can be exploited to improve efficiency via charge cooling, enabling higher compression ratios and potentially better volumetric efficiency. The hydroxyl (–OH) group and bound oxygen promote oxidation kinetics and yield higher laminar flame speed than gasoline under similar conditions. Conversely, the lower lower-heating value (LHV)—≈26.8 MJ/kg for ethanol vs ≈43 MJ/kg for gasoline—reduces heat release per unit mass. The stoichiometric air-to-fuel ratio of ethanol (≈9.0:1 by mass) implies a larger injected fuel mass for the same equivalence ratio compared with gasoline. Material compatibility must also be considered: ethanol can degrade certain elastomers and metals unless appropriate materials are used [8].

Hydrogen has gained strong relevance in recent years. Its simple molecular structure (H_2) and the absence of carbon make it appealing for deep decarbonization. Hydrogen can be used in fuel cells—predominantly Proton Exchange Membrane (PEM)—or burned directly in ICEs, although the latter requires specific adaptations. Its very wide flammability range, low minimum ignition energy, and high deflagration index increase the risk of backfire in the intake manifold and, less commonly, in the exhaust [9]. Hydrogen's laminar flame speed is significantly higher than gasoline's under comparable conditions, leading to faster combustion and typically reduced spark advance (closer to Top Dead Center, TDC) for optimal phasing. To control NO_x and mitigate combustion instabilities, ICE operation often uses very lean mixtures and/or high Exhaust Gas Recirculation (EGR) rates. In practice, global excess-air ratios $\lambda \approx 2–2.6$ are common compromises that balance performance, stability, and low NO_x emissions for spark-ignited H_2 engines, as illustrated in Figure 1 [10].

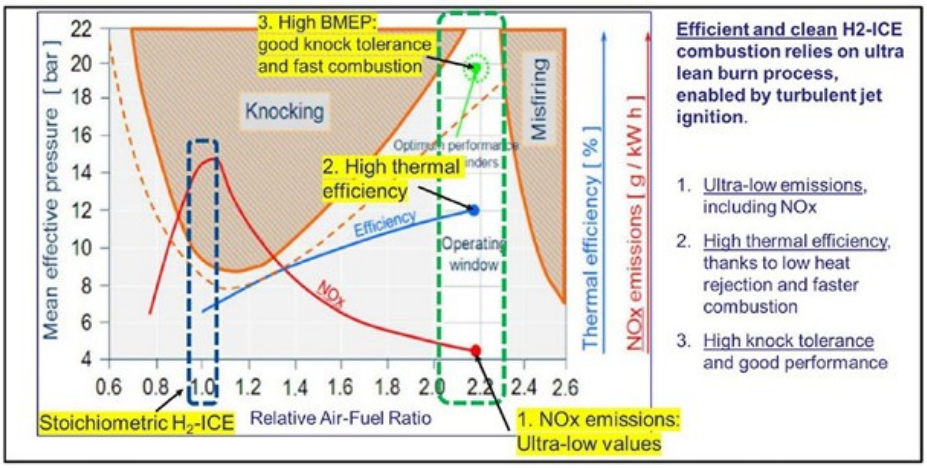

Figure 1: schematic combustion concepts for SI H2 ICE [10]

CHAPTER 1

1. E5 and E85 comparison in a single-cylinder, four-stroke, internal combustion engine

As noted in the introduction, ethanol is a well-known alternative fuel that can be used neat or blended with gasoline. However, its properties differ from conventional fossil fuels, so the impact on engine performance must be assessed.

This work aimed to analyze the use of E85 (85 vol% ethanol – 15 vol% gasoline) in a motorcycle engine, starting from a gasoline model calibrated with experimental data.

Engine characteristics are summarized in Table 1.

Engine Type	PFI-SI 4-stroke
Fuel Type	Gasoline (E5)
N. of Cylinders	1
Bore x Strone [mm]	96 x 64
Compression Ratio	12.5
N. of valves per cylinder	2 + 2
Exhaust Valve Opening	51
Intake Valve Opening	15
Exhaust Valve Closing	16
Intake Valve Closing	56
Air Metering	Naturally Aspired

Table 1: single-cylinder motorcycle engine

2. 1D model

The GT-SUITE 1D model is shown in Figure 2, and its correlation with experimental data is presented in Figure 3. The GT-Power fuel object used is Indolene, a GT-native representation of conventional gasoline.

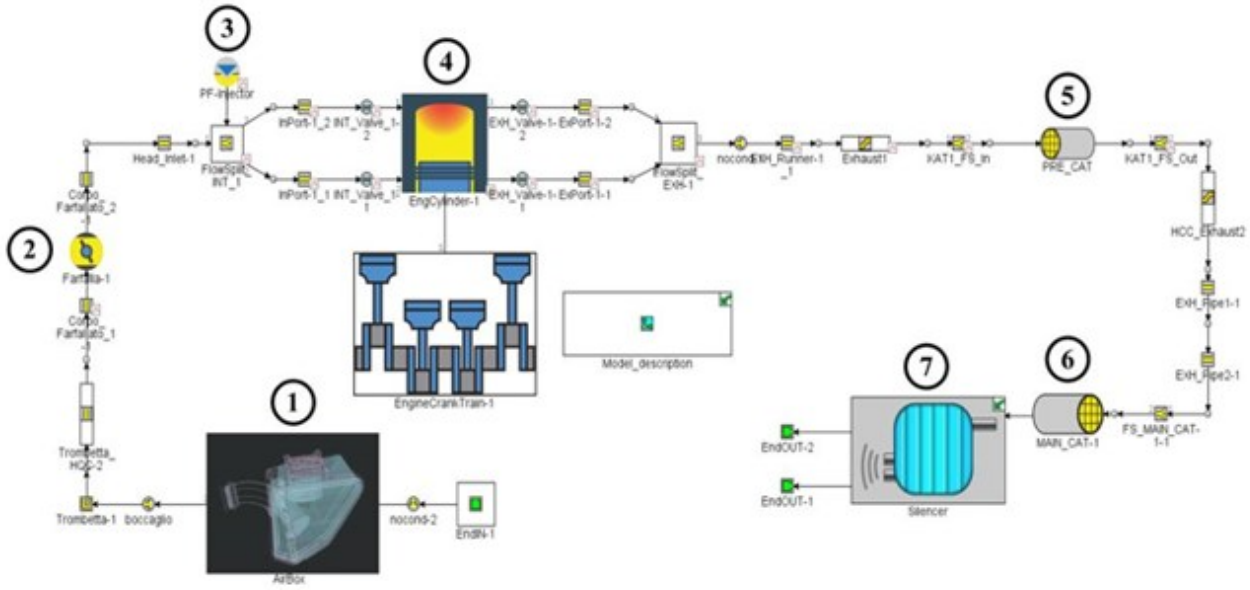

Figure 2: GT-Power model of the investigated engine

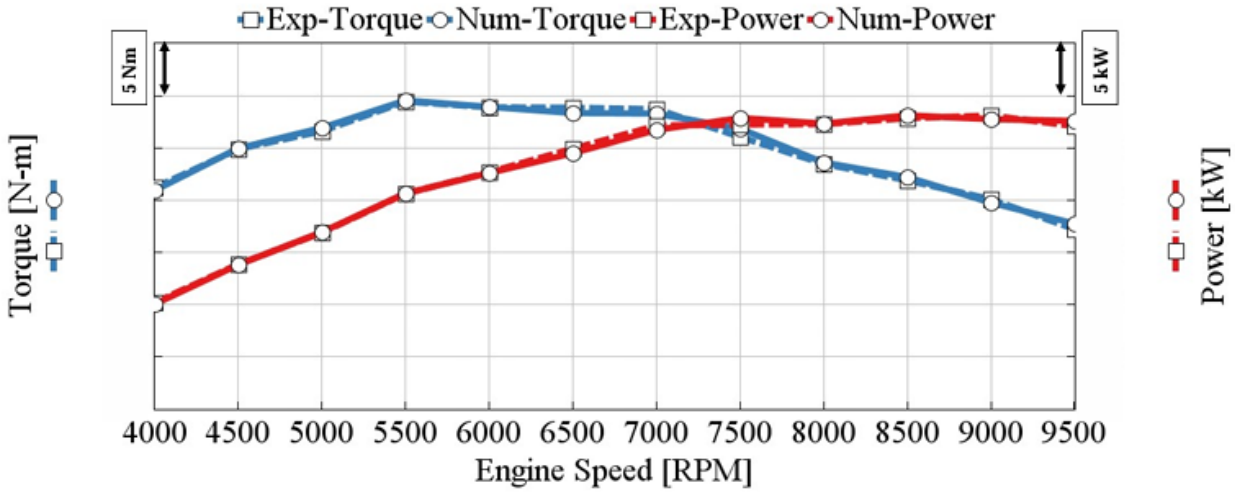

Figure 3: comparison between experimental and 1D numerical brake torque and power

Four engine operating points were considered (all at wide-open throttle, WOT):

- 4000 rpm — WOT
- 5000 rpm — WOT
- 7000 rpm — WOT
- 9500 rpm — WOT

Figure 4 shows good agreement between GT-Power simulations and experiments for cylinder absolute pressure at all operating points.

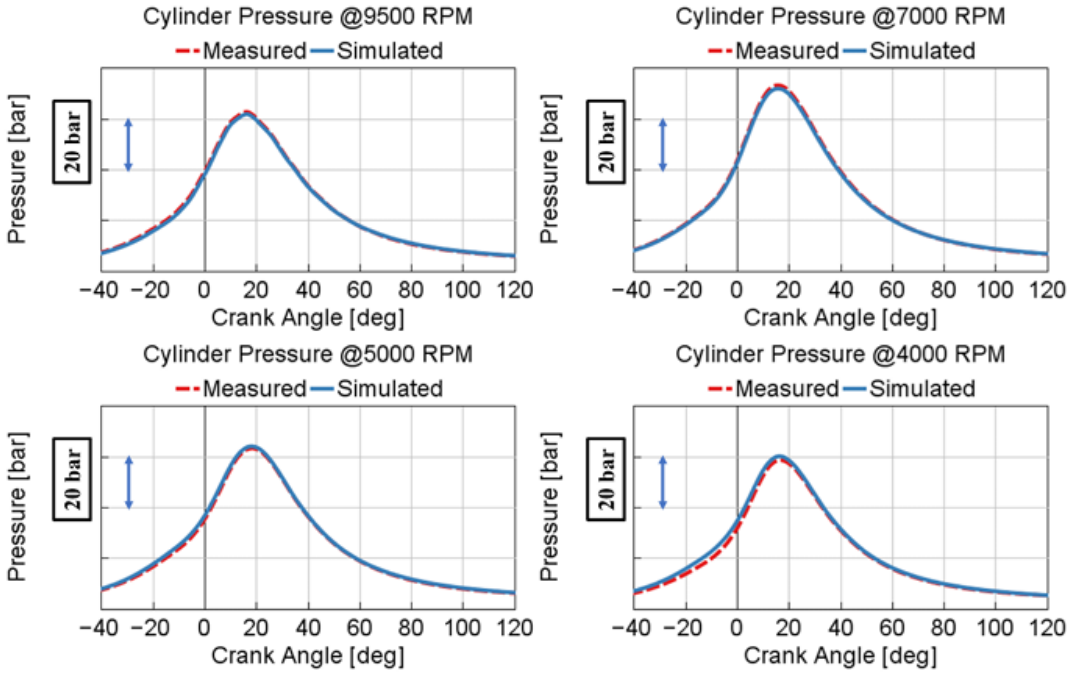

Figure 4: cylinder absolute pressure 1D simulation and experimental comparison

3. 3D model: model setup

The 3D CFD mesh was generated with the AVL FIRE M pre-configured workflow *Mesher*. This workflow ingests STL surface geometry; surfaces must be grouped and named according to the tool requirements [11]. Valve and piston motion, as well as local refinements, are handled automatically. By default, FIRE M creates two wall layers, and no modifications are allowed within this workflow manager. A mesh section in the valve region is illustrated in Figure 5.

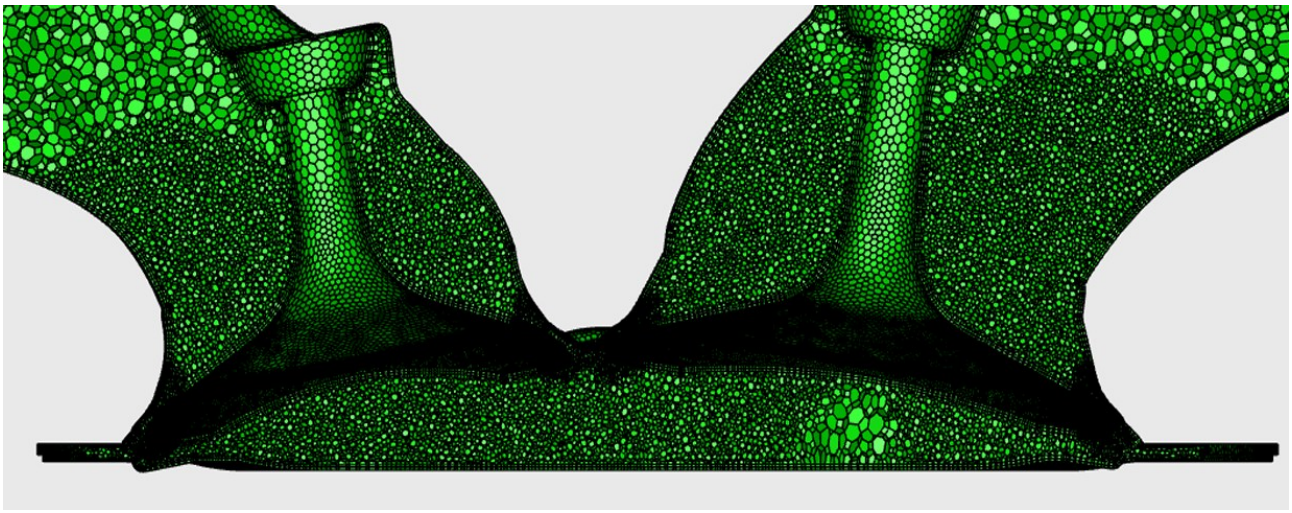

Figure 5: valve zone mesh refinement section

A correction that is commonly required when building a 3D CFD model—regardless of software—is the compression ratio. This has a non-negligible impact on results, especially for combustion simulations. The source of error is the discretized nature of the mesh versus the continuous real volume. In practice, the issue is corrected by moving STL surfaces (e.g., the piston crown) up or down so that the modeled compression ratio matches the target. For four-stroke engines such as this one, the maximum cylinder volume cannot be read directly from the mesh because the valves open. Therefore, the compression ratio must be computed from the firing TDC cylinder volume and the displacement volume (stroke × bore).

A simplified fuel surrogate was adopted in 3D by combining iso-octane (C_8H_{18}) and ethanol (C_2H_5OH). This differs from the Indolene fuel object used in GT-Power; however, the differences—especially during combustion—are mitigated by calibrating the combustion model. Conventional gasoline, referred to as E5, was modeled as 5 vol% ethanol + 95 vol% gasoline. Table 2 lists the main 3D-CFD models used.

Simulation type	RANS
Turbulence model	k-zeta-f
Wall heat transfer model	Standard
Near wall treatment	Hybrid
Combustion model	ECFM-3Z
Ignition model	Spherical
Laminar Flame Speed	Metghalchi & Keck

Table 2: 3D-CFD simulation main models

A premixed assumption was used for the engine cycle: the intake runner is filled with an air–fuel mixture, and the injection phase is omitted. This reduces computational cost because injections are not simulated and one cycle is sufficient to achieve cycle-to-cycle convergence.

The approximation is justified when the injected fuel fully evaporates before spark timing and the air–fuel mixing is well distributed in the combustion chamber. This was verified by running pre-cycle simulations to assess liquid-fuel distribution in the intake runner and cylinder. A schematic of the pre-cycle workflow is shown in Figure 6 [11]. For brevity, the detailed pre-cycle results are not reported here. Typically, ≈10 cycles are required for liquid-film convergence.

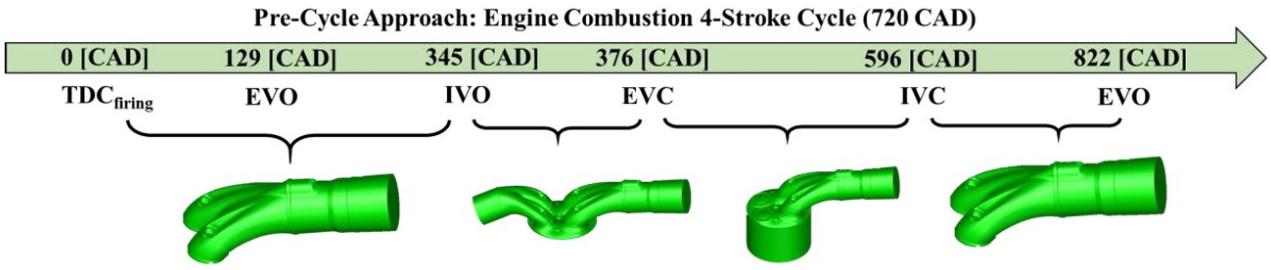

Figure 6: pre-cycle simulation approach

4. 3D model: E5 premixed simulation results

The premixed simulations were initialized with the exhaust and cylinder filled with CO_2, H_2O, and N_2 at mass fractions corresponding to ideal combustion products at stoichiometric air-to-fuel mass ratio. The intake runner was initialized with air and fuel only, at a prescribed mass ratio. Results in Figure 7 show good agreement between the 3D model and the 1D environment.

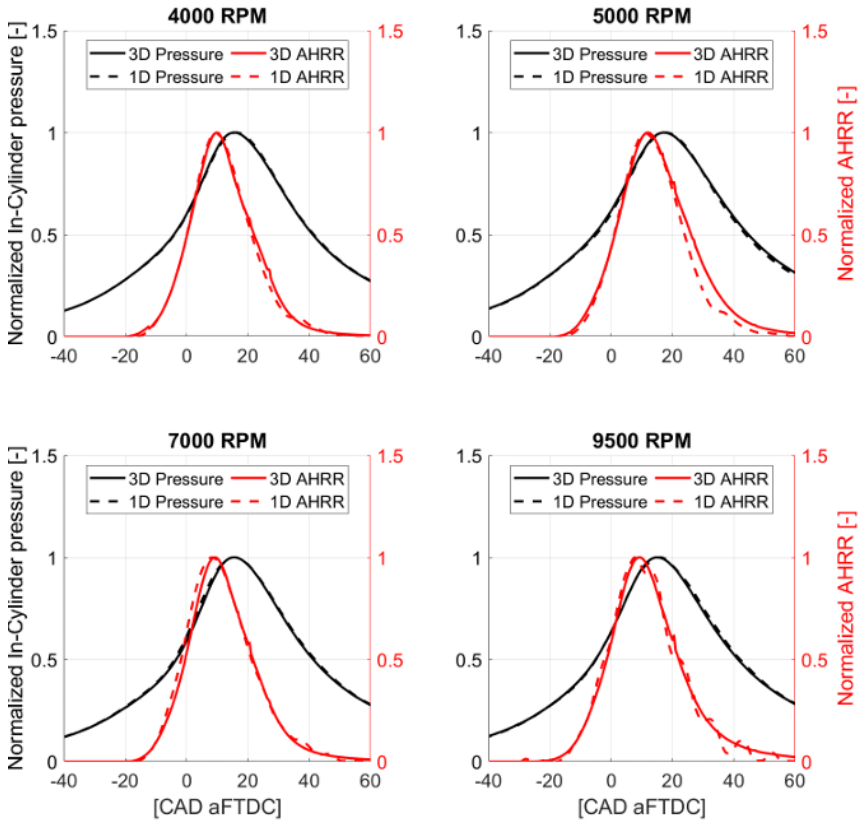

Figure 7: normalized 3D and 1D cylinder absolute pressure and AHRR comparison

The ECFM-3Z combustion model—key parameters are listed in Table 3—can be considered predictive for the phenomena investigated. These parameters were held constant across all operating points:

- Stretch factor: a multiplier for turbulence–combustion interaction; higher values imply higher turbulent combustion speed
- Initial Flame Surface Density: initial surface-to-volume ratio of the flame front; it represents the initial energy provided by the spark plug
- Mixing-model parameter: related to diffusive combustion; because the mixture is already premixed here, its influence is negligible

Strech factor [-]	0.69
Initial Flame Surface Density [1/m]	1250
Mixing model parameter [-]	1

Table 3: ECFM-3Z combustion parameters

5. 3D model: E85 model setup

After validating the gasoline cases, the fuel object was replaced with E85 and the E85 analysis was performed at 9500 rpm, WOT.

An optimization of injected fuel mass was carried out in the GT-Power environment using maximum IMEP as the constraint. The equivalence ratio was adjusted from 1.186 (E5) to 0.93 (E85). This provided a new, self-consistent set of boundary conditions for the 3D simulations.

To account for the higher latent heat of evaporation of E85 relative to E5, a new pre-cycle was run to evaluate liquid-fuel evaporation and distribution in the runner and cylinder. Any liquid film remaining after IVC was assessed to ensure it would not affect volumetric efficiency.

Because the premixed approach is used, latent-heat effects are not directly captured during the cycle. Therefore, we introduced a temperature-decrease reinitialization based on the following energy balance (Equation 1), where ΔT is the air temperature decrease, $c_{p,air}$ is the constant-pressure specific heat of air, m_{air} is the in-cylinder air mass trapped for the E85 case, m_{E85} and m_{E5} are the trapped fuel masses for E85 and E5, Y_{E85} and Y_{E5} are the fractions of fuel effectively participating in air-temperature decrease, and h_{E85} and h_{E5} are the latent heats of evaporation:

$$m_{air} c_{p,air} \Delta T = (m_{E85} h_{E85} Y_{E85} - m_{E5} h_{E5} Y_{E5})$$

Equation 1: air temperature decrease relationship

Rearranging gives Equation 2, expressed via the equivalence ratios ϕ_{E85} and ϕ_{E5}, and the stoichiometric air-to-fuel mass ratios $(A/F)_{s,E85}$ and $(A/F)_{s,E5}$:

$$\Delta T = \left[\frac{\Phi_{E85} h_{E85}}{\left(\frac{A}{F}\right)_{s,E85}} Y_{E85} - \frac{\Phi_{E5} h_{E5}}{\left(\frac{A}{F}\right)_{s,E5}} Y_{E5} \right] * \frac{1}{c_{p,air}}$$

Equation 2: air temperature decrease

A sensitivity analysis was performed by sweeping equivalence ratio and spark advance (SA):

- Φ: 1.186, 1.1, 1.0, 0.95, 0.9, 0.8
- SA: −50, −45, −40, −35, −30, −27, −25, −15 CAD

For each equivalence ratio, a spark-advance sweep was conducted. Reference values for gasoline (E5) are $\Phi = 1.186$ and SA = −25 CAD. Combustion-model parameters from the gasoline correlation were retained.

6. 3D model: E85 premixed simulation results

For each simulated case, gIMEP* and combustion efficiency were evaluated. Here, gIMEP* denotes the *gross* indicated mean effective pressure, computed from 651 to 800 CAD. Combustion efficiency is defined by Equation 3:

$$Combustion\ efficiency = \frac{Q_c}{m_{f,tot} * LHV}$$

Equation 3: combustion efficiency

where $Q_c$is the total released heat, $m_{f,\text{tot}}$is the total fuel mass, and LHV is the lower heating value.

Figure 8-Figure 9 report gIMEP* and **c**ombustion efficiency for all analyzed cases; the gasoline baseline is marked with yellow stars.

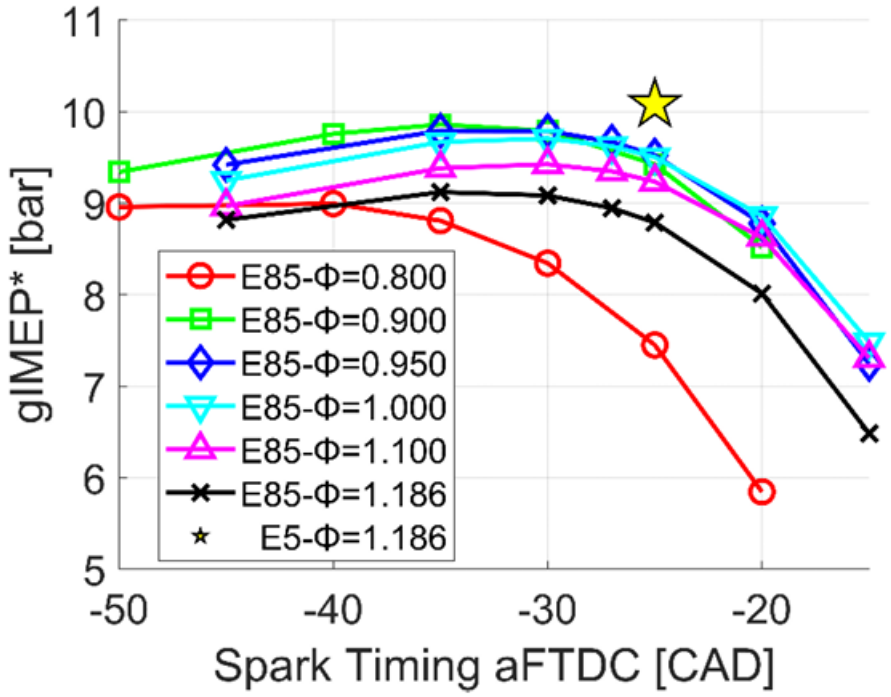

Figure 8: gIMEP vs Spark Advance at each equivalence ratio*

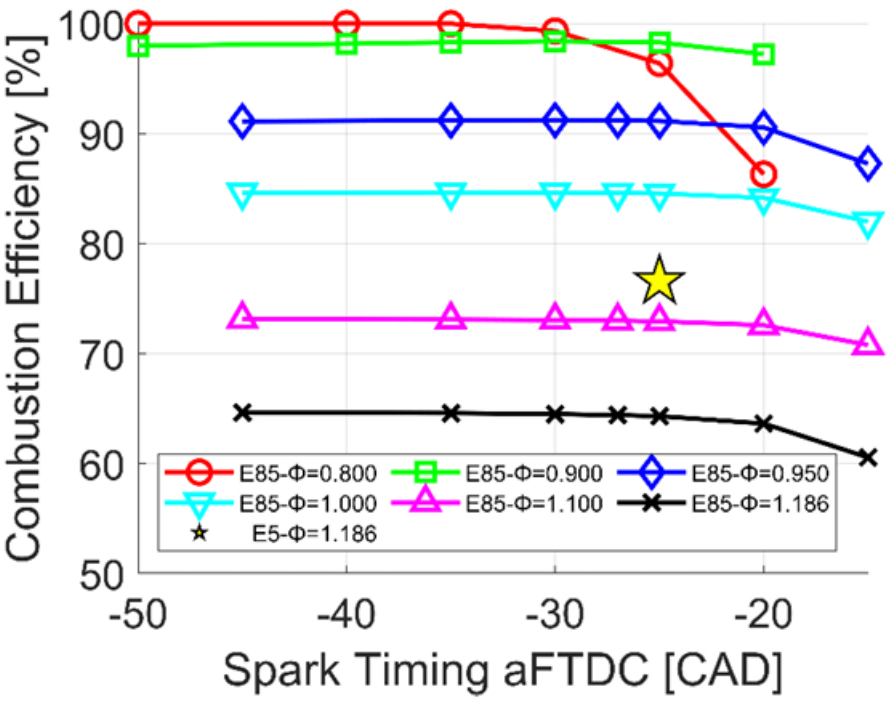

Figure 9: combustion efficiency vs Spark Advance at each equivalence ratio

Based on these results, the best operating window is ϕ between $[0.90, 0.95]$ with SA ≈ −30° to −25° aFTDC.

It is noteworthy that E85 delivers higher combustion efficiency—up to ~22% above E5 when $\phi < 1$. This arises because, at lower equivalence ratios, the denominator of Equation 3 decreases. Overall performance (gIMEP*) is lower than E5; however, for ϕ slightly below 1 and SA = −30°, gIMEP* is only ~2.2% below the E5 case.

Figure 10 shows the turbulent burn duration (10–90 CAD). Owing to ethanol's higher laminar flame speed, combustion is faster. This capability can be leveraged to operate at leaner mixtures while retaining acceptable performance—unusual for traditional gasoline engines, which typically require at least stoichiometric operation. Moreover, because E85 requires a larger injected fuel mass than E5 at the same equivalence ratio and has more than twice gasoline's latent heat of evaporation, the knock risk is strongly mitigated.

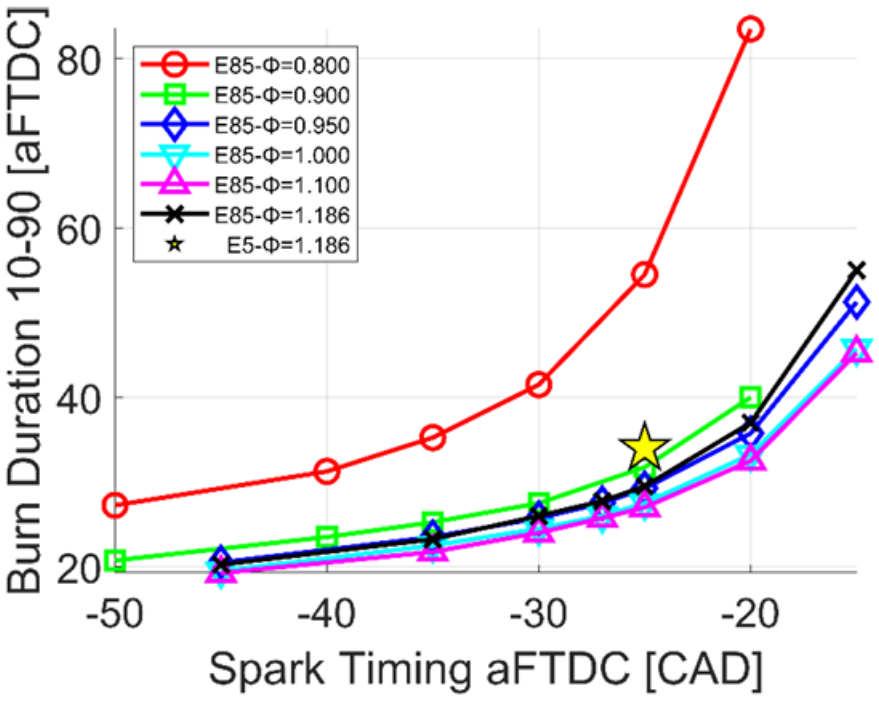

Figure 10: turbulent combustion duration (10-90 CAD)

7. Conclusions

This study examined E85 in a single-cylinder, four-stroke motorcycle ICE. Starting from a gasoline model calibrated on experimental data, the fuel was replaced with E85 in both the 1D and 3D models.

A sweep over equivalence ratio and spark advance was performed, and results were compared with the E5 case at the same operating point (9500 rpm, WOT). The analysis indicates that both ϕ and SA must be retuned for E85.

Lower equivalence ratios yield higher efficiency at a given spark advance. The optimal range is $\phi = 0.90–0.95$, consistent with the GT-Power optimization (optimal $\phi = 0.93$). Detonation risk is reduced thanks to the higher injected mass with E85 (at fixed ϕ) and ethanol's significantly larger latent heat of evaporation.

From this perspective, E85 is a promising replacement for conventional gasoline (E5): it can deliver greater efficiency with only a modest performance penalty after appropriate calibration.

CHAPTER 2

1. E85 combustion with detailed chemical-kinetics mechanism in a four-cylinder, four-stroke, Port Fuel Injection motorcycle ICE

During my studies, I investigated E85 and E5 combustion using AVL FIRE M with detailed chemical-kinetics mechanisms. The approach is similar to the previous chapter: injection is replaced by a premixed setup in which the intake runner is filled with air and fuel to achieve a target equivalence ratio.

The engine under study is the UNIMORE Formula four-cylinder, four-stroke ICE with:

- Compression ratio: 12.5
- Stroke: 46 mm
- Bore: 70 mm
- Connecting rod length: 96.35 mm
- Number of valves (intake + exhaust): 2 + 2

This engine is widely used in our department, and experimental data (e.g., brake power) are available for both gasoline and E85. Figure 11 shows the geometry used for the 3D CFD analyses. Only one cylinder is considered here.

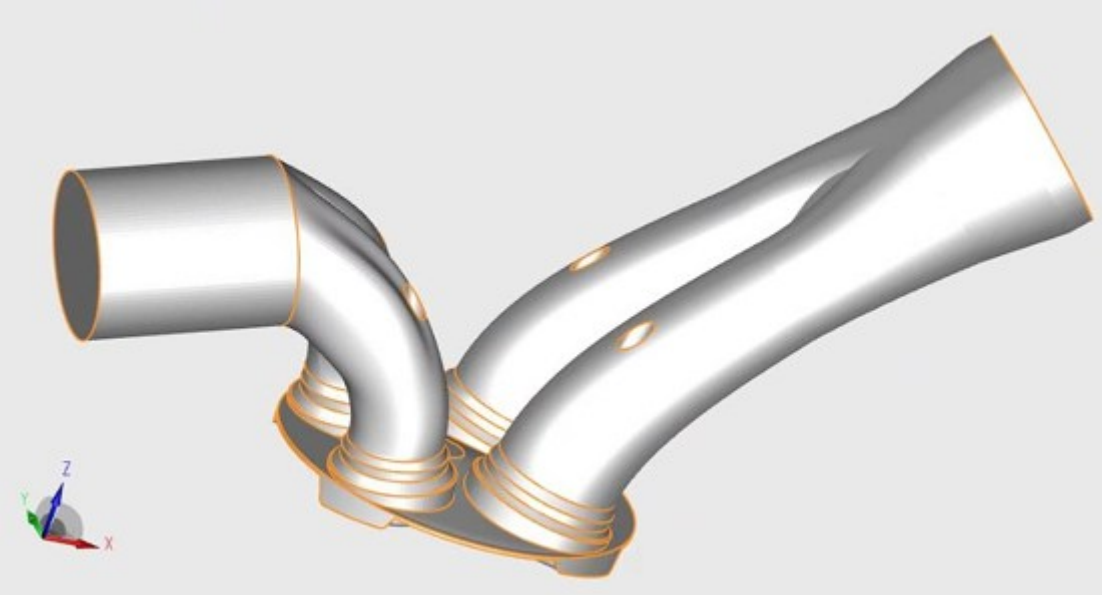

Figure 11: UNIMORE Formula Student ICE

GT-Power models, ready to use, were already available and calibrated with experimental data from the UNIMORE test bench. Recent SAE regulations allow the use of E85 and attribute an additional “efficiency score” when E85 replaces gasoline.

The goal was to assess combustion modeled with detailed chemistry (DC) instead of the ECFM-3Z approach used earlier. DC relies on sets of species and reactions validated against experiments (e.g., perfectly stirred reactors). In DC, a dedicated Ordinary Differential Equation (ODE**)** is solved within the CFD domain for each specie (unless specific cost-reduction strategies are applied). In practice, 50–150 species are typical.

In this work, we adopted the mechanism by Yunchao Wu et al. [12], which contains 165 species and 839 reactions. It derives from a reduced LLNL mechanism and is tailored for ICE use. Validity ranges are:

- Equivalence ratio: 0.5–1.3
- Absolute pressure: 1–100 atm
- Initial temperature: 750–1800 K (includes low-temperature chemistry)

2. Model setup

Since RANS turbulence model (k-zeta-f) was used, a symmetry plane was exploited so that half geometry could be simulated, reducing computational effort. To obtain representative fields of flow, pressure, and species, a cold-flow simulation was first run from 130 CAD (before IVO) to 680 CAD (before SOC), applying boundary conditions from the GT-Power 1D model—consistent with the approach described earlier.

The General Gas Phase Reaction (GGPR) module must be activated to use DC mechanisms. A new material was created in FIRE M containing all mechanism species and properties. However, for cold-flow runs aimed at developing a mature flow before SOC, using the complete mechanism is costly. To limit runtime, a “light” material was defined from the original mechanism, retaining only C_2H_5OH, iC_8H_{18}, O_2, N_2, H_2O, CO_2; transport properties (e.g., viscosity, thermal conductivity) were kept identical to the full mechanism. Nevertheless, E5 and E85 surrogates in combustion runs follow the same component split as before (iC_8H_{18} and C_2H_5OH).

Ignition is triggered by heat addition within a spherical region whose cells are identified by a sphere object. FIRE M’s ignition model includes:

- Coordinates: spark-center position
- Spark Timing [deg]: ignition start
- Flame Kernel Size [mm]: sphere radius (default 3 mm)
- Ignition Duration [s]: model action time (default 3e-4 s)
- Energy Factor [–]: calibration coefficient (higher → stronger ignition)

Additional multipliers control turbulence–chemistry interaction and heat release computed from kinetics. As highlighted in [13], detailed chemistry can deliver predictive results across regimes (e.g., thickened flame / flamelet) when parameters are calibrated to experiments.

3. Simulation results: E85

A series of combustion simulations was performed starting from the 680 CAD cold-flow state. The Energy Factor served as calibration parameter, from the default value 1 up to ≈50. Figure 12 reports cylinder absolute pressure for EF = 1, 20, 30, 38, 43. The label “NM” denotes *new material* (complete mechanism), while “EF” indicates *Energy Factor*. Cases not plotted up to 750 CAD were stopped due to poor matching with the target.

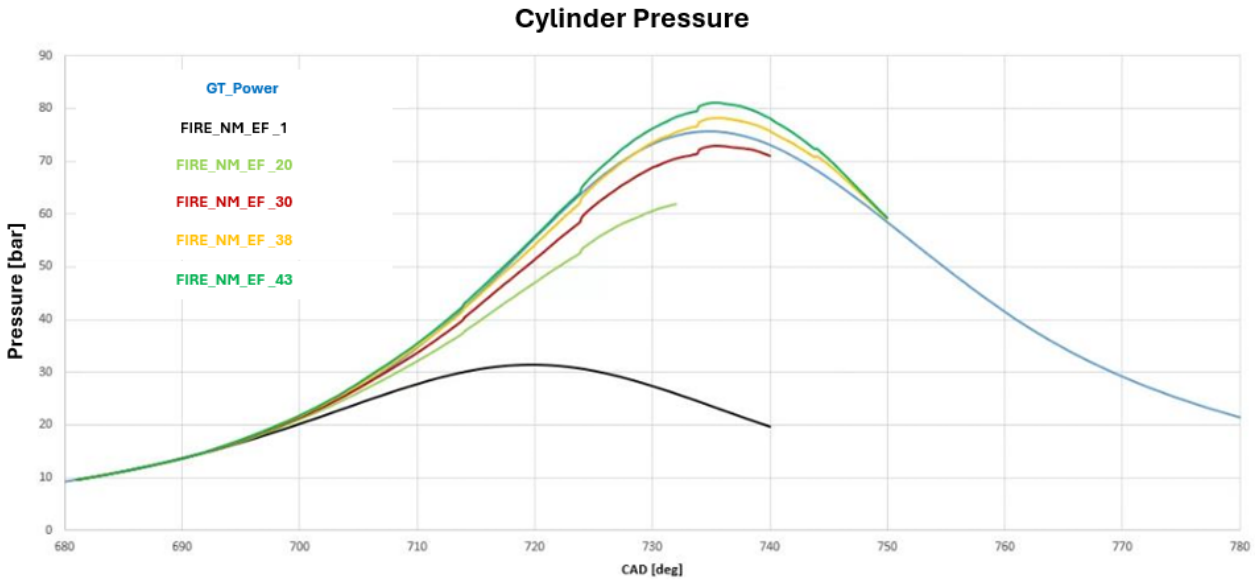

Figure 12: cylinder absolute pressure at different Energy Factor – E85 fuel

The GT-Power reference (blue curve) is used for comparison. From the AHRR in Figure 13, EF = 1 does not ignite the mixture; values around ~40 yielded the best correlation.

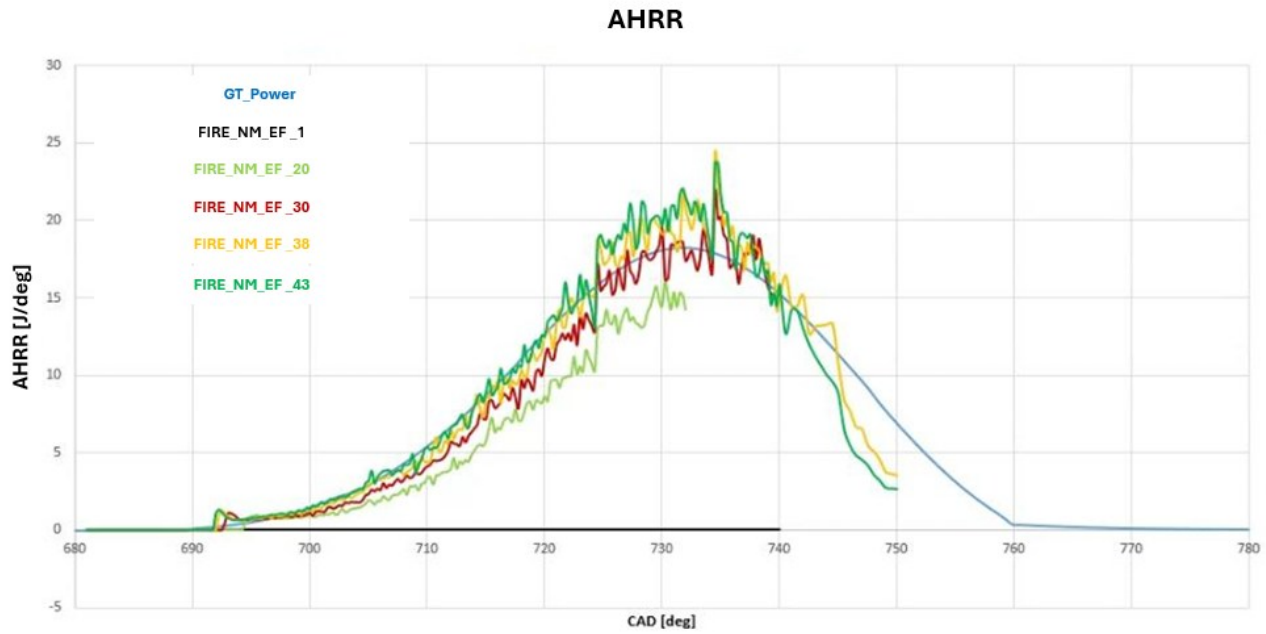

Figure 13: Apparent Heat Release Rate (AHRR) [J/deg] – E85 fuel

Figure 14 and Figure 15 show Cumulative Heat Released (CHR) and Mass Fuel Burnt (MFB), respectively. Simulated CHR is lower than the target, whereas MFB approaches 100%, indicating that nearly all trapped E85 is consumed. The ratio between simulated CHR and the CHR estimated as (E85 LHV × mass burnt) is ≈0.87. This gap is attributable to DC predicting intermediate species (e.g., H_2, CO) that retain chemical energy and thus reduce the released-heat total.

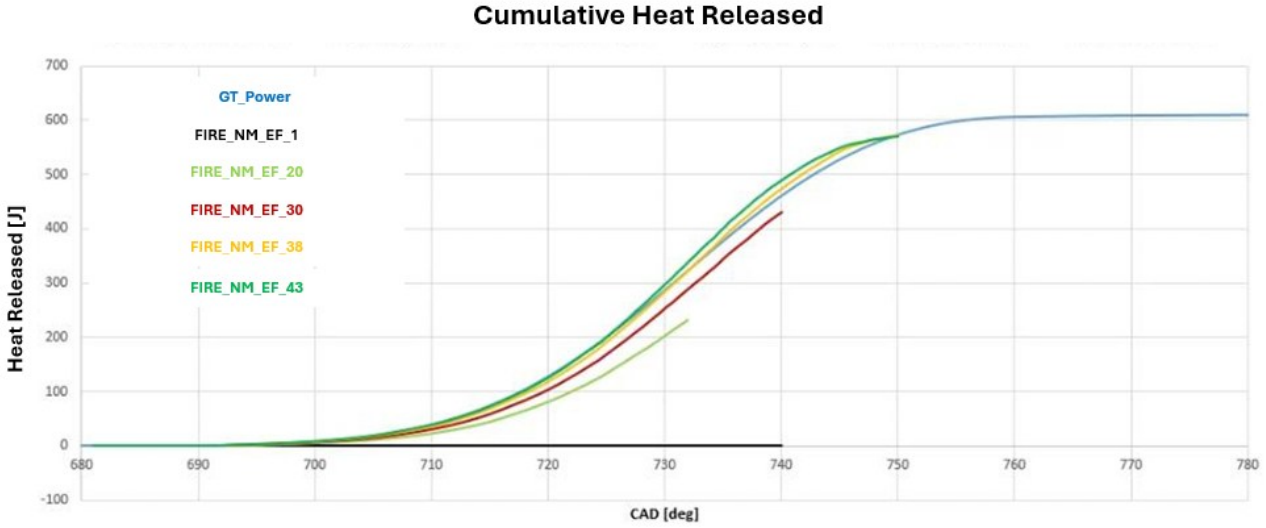

Figure 14: cumulative heat released [J] – E85 fuel

The simulated IMEP (computed from 680 to 780 CAD) exceeds the GT-Power value by 4.3%, which is acceptable for our purposes. Mesh “rezone” cycles—required by FIRE M’s mesh-motion strategy—can introduce discontinuities: vertex motion deforms cells, so remeshing is needed and fields are mapped between meshes. Interpolation across differing cell shapes may cause visible steps; this can be mitigated by reducing the angle range between consecutive sets so that cell shape/volume remain similar.

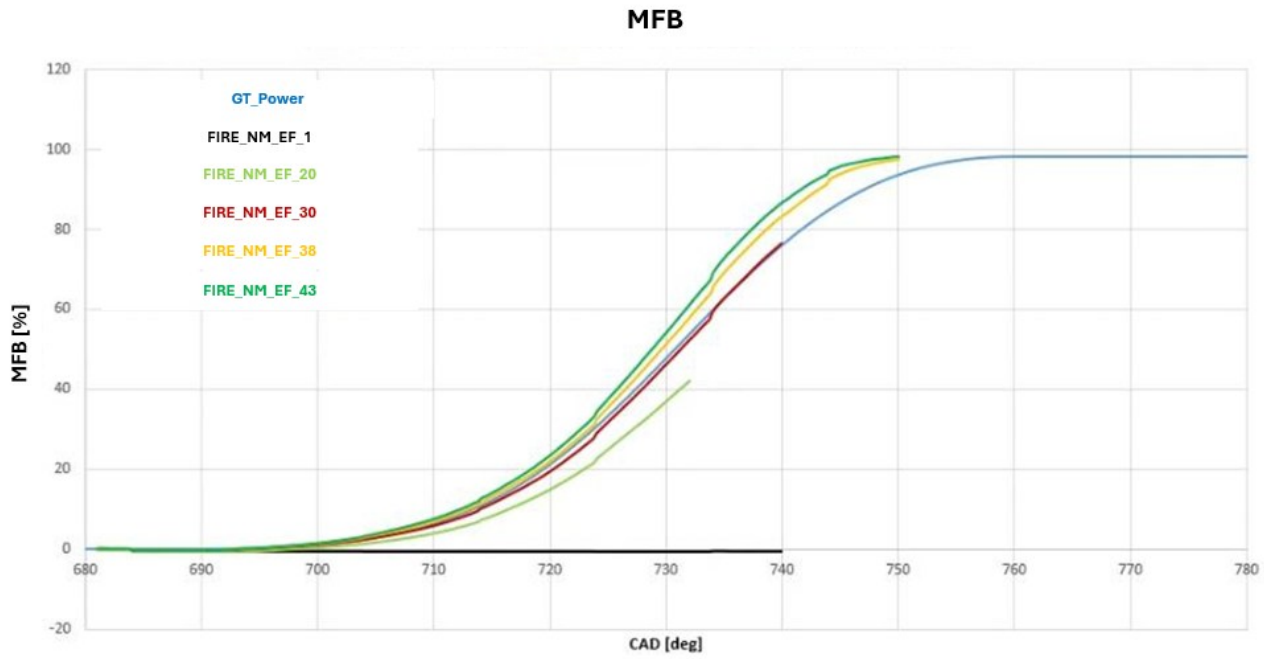

Figure 15: Mass Fuel Burnt [%] – E85 fuel

4. Simulation results: E5

This section presents the E5 combustion results. The same methodology as E85 was used, together with the same detailed-chemistry mechanism and combustion-model parameters.

It was observed that, for E5, Energy Factor calibration alone was insufficient to obtain a proper match. As Figure 16 shows, even with EF = 55 the simulated cylinder absolute pressure remains below the reference. Further increasing EF does not close the gap and, in addition, produces unphysical temperatures in the spark region (≈5000 K), which is unacceptable and may cause numerical divergence.

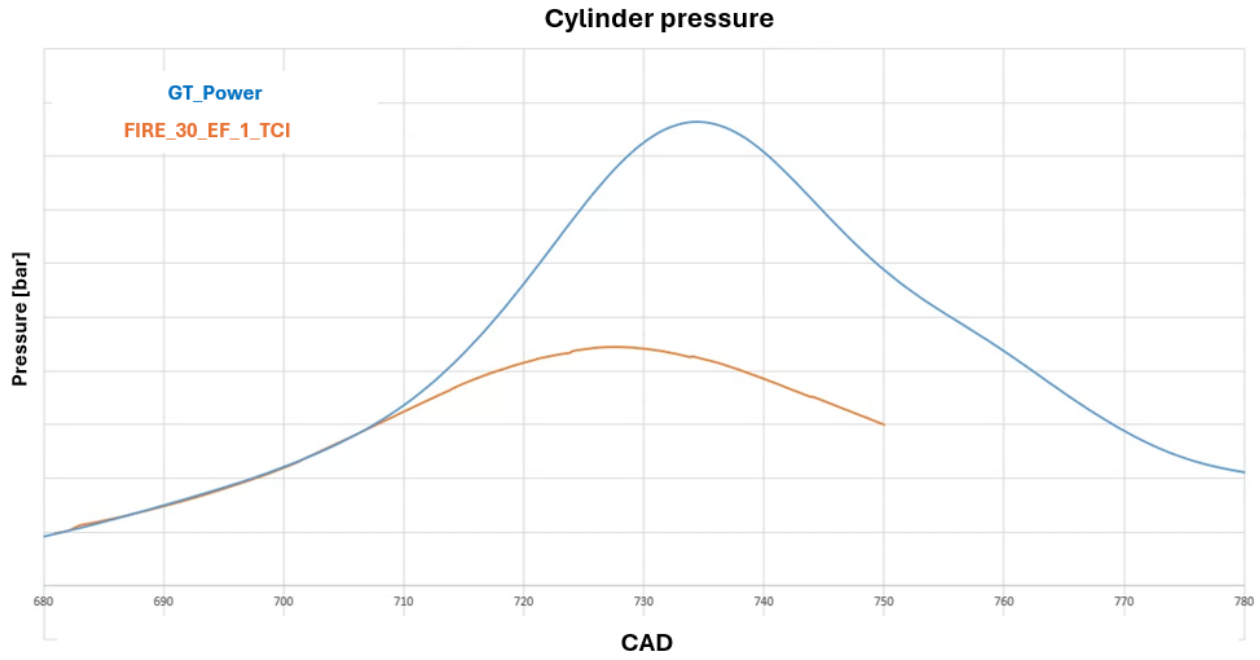

Figure 16: cylinder absolute pressure – Energy Factor: 55; TCI: 1 – E5 fuel

A practical remedy is to increase turbulence chemistry interaction (TCI) intensity. As noted earlier, this is controlled by a global multiplier on TCI. In this case, setting TCI ≈ 1.3–1.35 while keeping EF = 38 (the same as the E85 case) yields good agreement with experiments. This behavior can be ascribed to the different combustion regimes of E5 and E85.

Figure 17Figure 18 report cylinder absolute pressure and AHRR at different TCI values. With TCI = 1.35 and EF = 38, the 3D simulated IMEP (computed from 680 to 780 CAD) is 2.7% lower than the target, which demonstrates good agreement with the experimental reference.

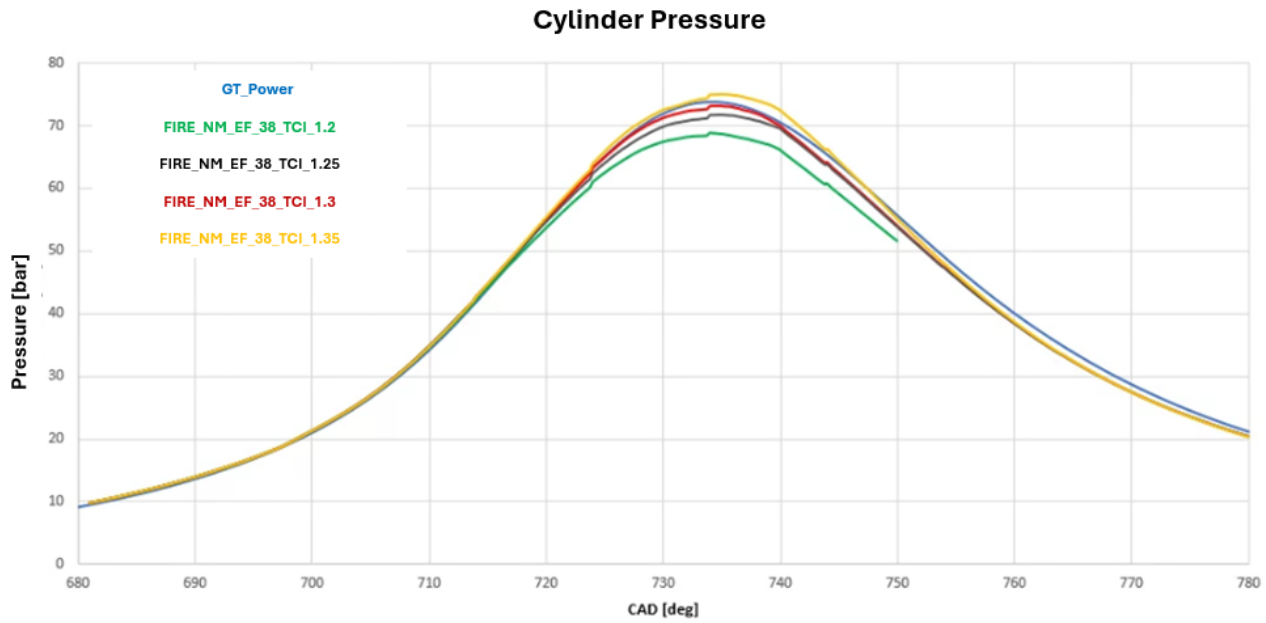

Figure 17: cylinder absolute pressure [bar] – E5 fuel

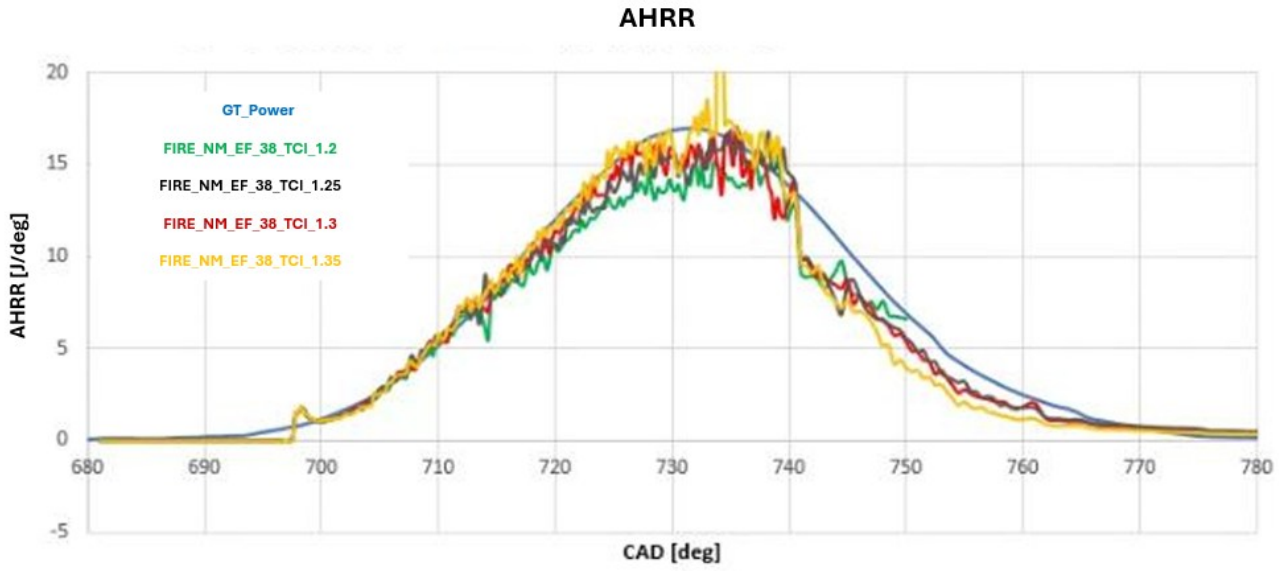

Figure 18: apparent heat release rate [J/deg] – E5 fuel

5. E85 and E5 results comparison

A comparison of E85 and E5 combustion results is presented. From Cumulative Heat Released (CHR) (Figure 19), E85 shows higher values, primarily because the trapped fuel mass is slightly larger than for E5. Moreover, E5 exhibits poorer energy conversion: the ratio between simulated CHR and the theoretical heat (LHV × mass burnt) is 0.87 for E85 and 0.80 for E5, implying that E5 products retain more unoxidized chemical energy.

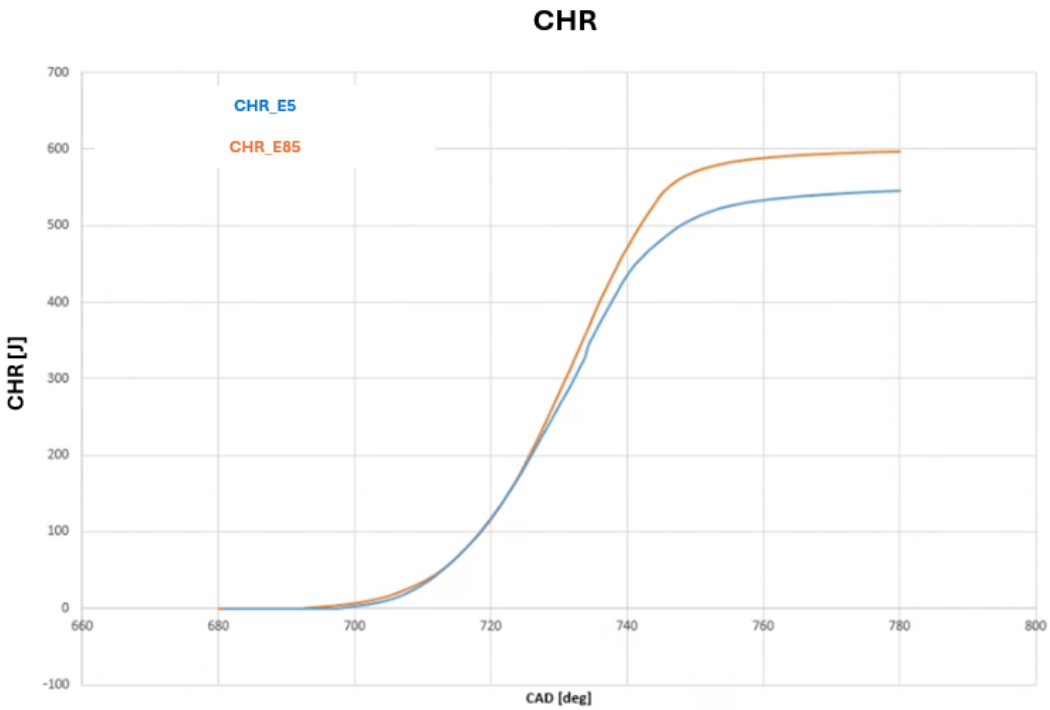

Figure 19: CHR comparison [J]

The turbulent burn duration (MFB10–MFB90) is 30.5 CAD for E85 and 32.5 CAD for E5. The higher ethanol content in E85 increases laminar flame speed, shortening burn duration, even though the equivalence ratio is slightly above unity in both cases. The oxygen bound in ethanol acts as a premixed oxidizer that facilitates oxidation processes [14].

Regarding pollutants, E85 yields lower emissions, as shown in Figure 20–Figure 22. The higher CO_2 for E85 should be attributed to more complete combustion, and for equal ϕ the injected fuel mass is higher with E85 due to its lower stoichiometric air–fuel ratio.

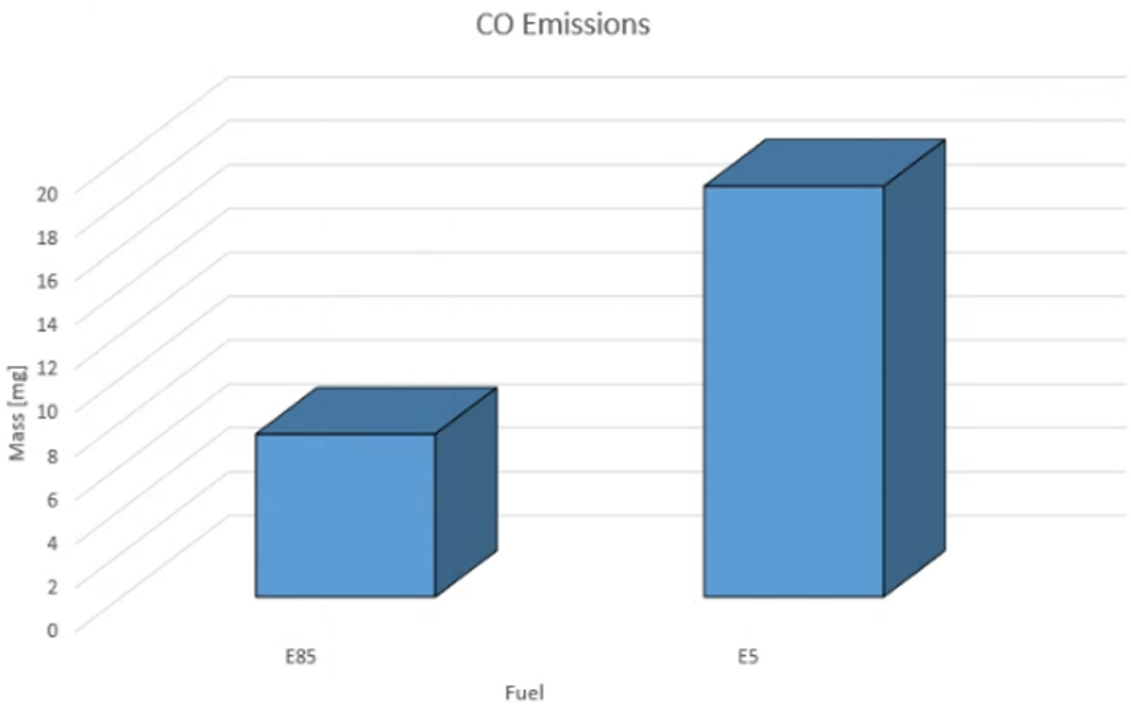

Figure 20: CO mass comparison

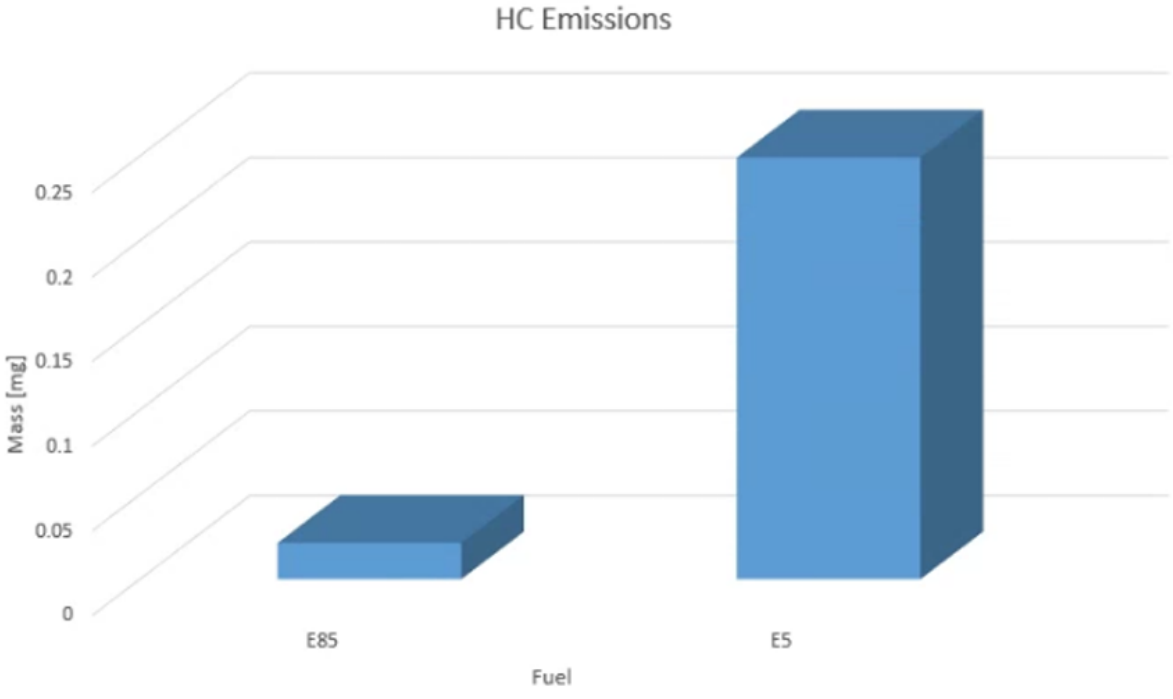

Figure 21: HC mass comparison

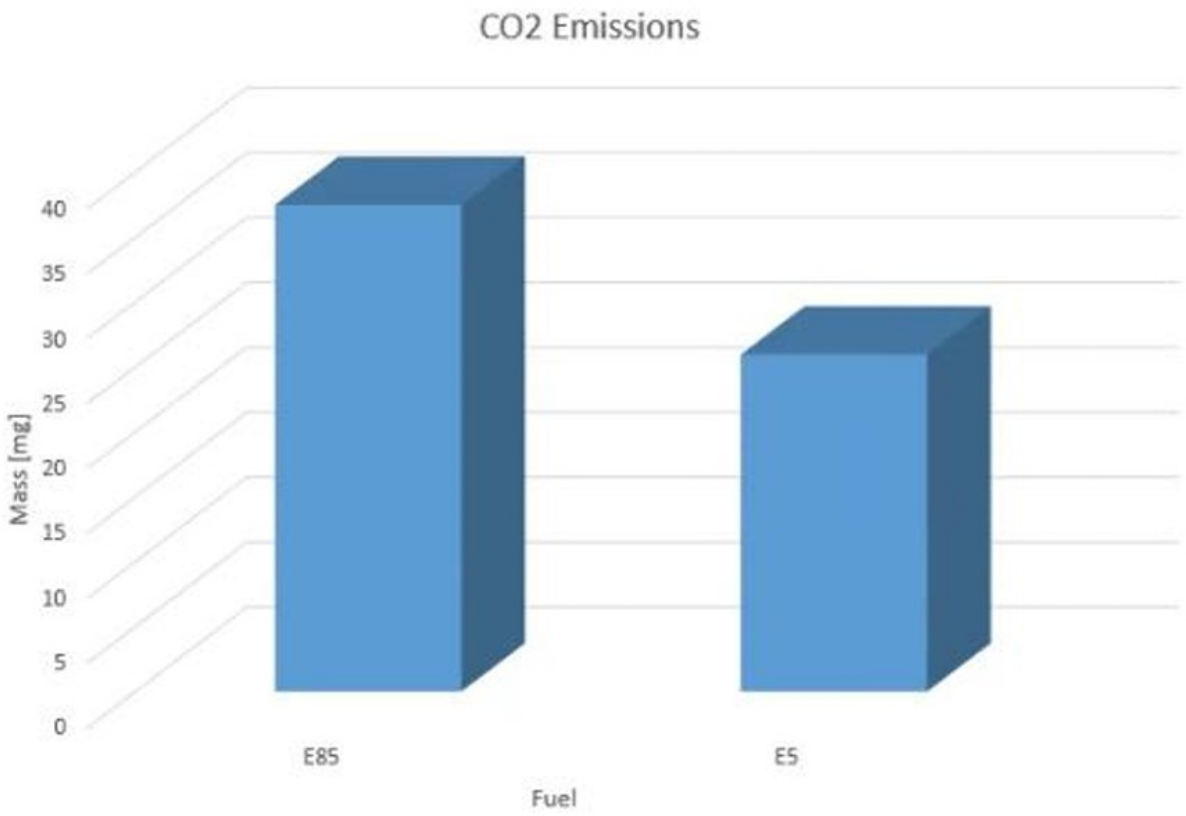

Figure 22: CO_2 mass comparison

The 3D simulated IMEP (680–780 CAD) is 11.12 bar for E85 and 10 bar for E5. Spark timing is 690 CAD for E85 and 695 CAD for E5. Together with E85's higher laminar flame speed, this leads to higher temperatures and thus higher thermodynamic efficiency, translating into higher IMEP.

Overall, the emissions comparison indicates cleaner combustion with E85 than with E5: CO and HC are lower, whereas CO_2 is higher, consistent with more complete oxidation.

6. Conclusions

This study compared the use of E85 and E5 in a four-cylinder**,** four-stroke PFI motorcycle ICE using AVL FIRE M with Detailed Chemical Kinetics.

Experimentally calibrated GT-Power 1D models provided boundary conditions for the 3D CFD simulations. The General Gas Phase Reaction module was enabled, and a new material object was built containing all mechanism species and properties.

Cold-flow simulations adopted a premixed approach and used a light material (C_2H_5OH, iC_8H_{18}, O_2, N_2, H_2O, CO_2) to reduce computational time. Subsequent combustion simulations employed the complete mechanism. Two parameters were used for calibration against experimental data: Energy Factor (EF) and Turbulence–Chemistry Interaction (TCI). For E85, EF = 38 provided adequate correlation; for E5, TCI ≈ 1.3 was required (with EF = 38) to achieve similar agreement. The need for different settings likely arises from different combustion regimes for E85 vs. E5.

In summary, E85 delivers higher thermodynamic efficiency (higher IMEP), cleaner emissions (lower CO and HC, higher CO_2), and shorter burn duration than E5 under the tested conditions, at the cost of a higher injected fuel mass due to its chemistry and stoichiometry.

CHAPTER 3

Opposed-piston ICE fueled by hydrogen — scavenging and combustion analysis

OpenFOAM is a widely used open-source platform for 3D CFD. It supports a broad set of applications—external aerodynamics, detailed combustion, internal flows, etc. Being open, it allows users to extend features, implement new models, and run diverse simulation strategies. The trade-off is the lack of a native GUI and a steeper learning curve requiring solid UNIX skills.

I began using OpenFOAM self-taught during my Ph.D. and collaborated with Tommaso Lucchini's research group at Politecnico di Milano to study unconventional ICE architectures, specifically the opposed-piston (OP) engine. OP engines are not new—production dates back to the early 1900s in Germany, the USA, the UK, France, and Russia [15], with extensive use during and after WWII.

The novelty here is the use of hydrogen. Burning H_2 directly in ICEs can provide mechanical power while leveraging OP uniflow scavenging to enhance delivery and trapping ratios and keep exhaust residuals low. As noted in the introduction, hydrogen demands careful control of the combustion process; models originally developed for carbon-based fuels do not fully capture all H_2-specific behaviors [16].

The Politecnico di Milano group developed libraries tailored to ICEs, known as LibICE. Among them are validated hydrogen-combustion models that account for thermo-diffusive instabilities, based on a modified Weller combustion model [17] and dedicated mesh-motion strategies, including OP motion.

In this work, a single-cylinder OP engine with a displacement of 0.905 L (≈30 kW at 3000 rpm) is simulated in OpenFOAM 8 (ORG version) to investigate scavenging and combustion. Table 4 summarizes the main engine properties.

N. of cylinders	#	1
N. piston per cylinder	#	2
Bore	mm	82
Stroke	mm	90
Connecting Road Length	mm	160
Displacement	cm3	950.6
Geometric Compression Ratio	-	14
Crankshaft offset	deg	10

Table 4: Opposed Piston characteristics

A 1D GT-Power model was built by the DIEngineLAB group to provide boundary conditions for the 3D CFD.

3D CFD workflow summary can be subdivided as follows:

- Generate high-quality static volume meshes using OpenFOAM's native blockMesh and snappyHexMesh.
- Run cold-flow cases with LibICE (solver coldEngineDyMFoam) to compute gas-exchange performance using RANS turbulence (k-epsilon) approach.
- Run combustion cases at different spark timings, imposing uniform air–fuel mixing using RANS turbulence (k-epsilon) approach.

1. Mesh generation

OpenFOAM meshes were created with two tools: blockMesh and snappyHexMesh.

blockMesh produces structured hexahedral meshes with precise control of vertices, edges, and faces—ideal for medium-complexity components. It is less suitable for highly complex geometries typical of ICE ducts. snappyHexMesh excels at complex shapes: it starts from a uniform, structured background mesh (usually a parallelepiped of hexahedra generated by blockMesh) and then snaps/refines to STL surfaces [18].

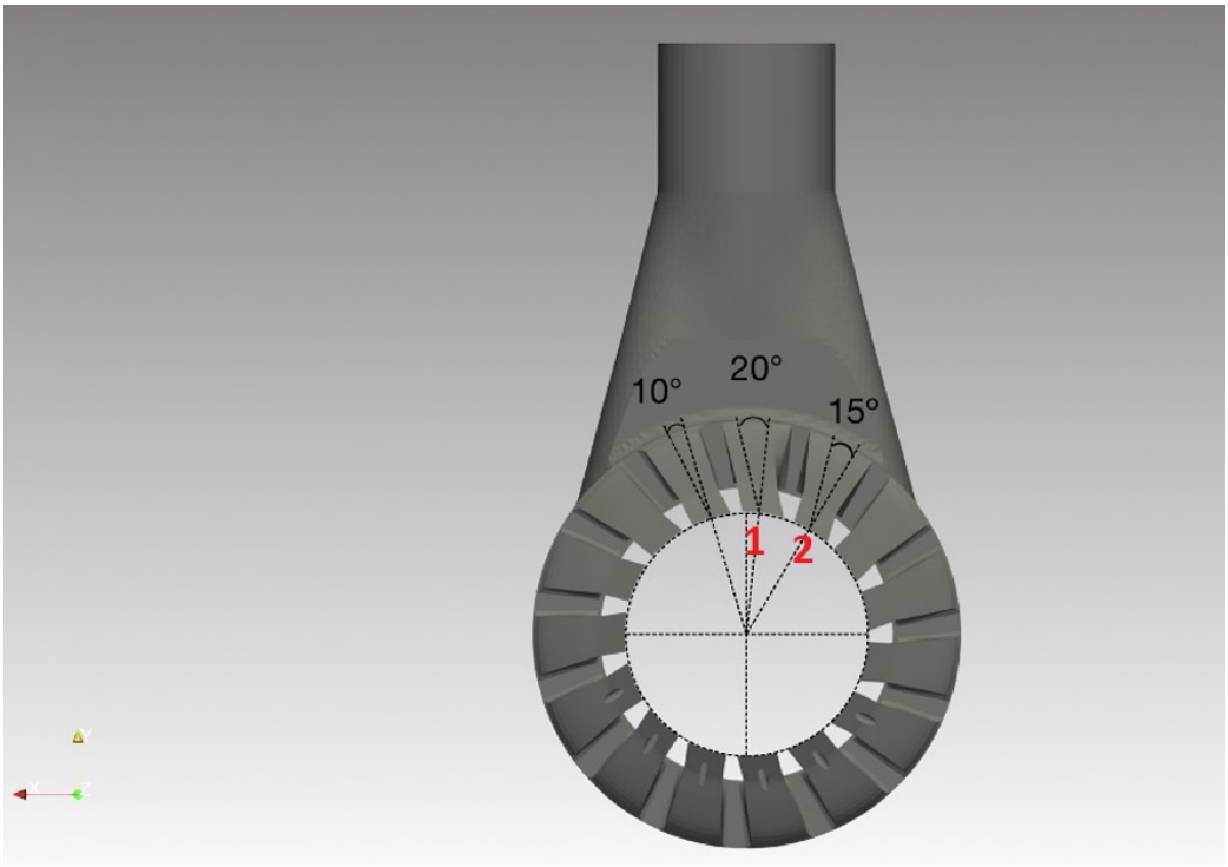

Figure 23: Opposed Piston Intake Geometry

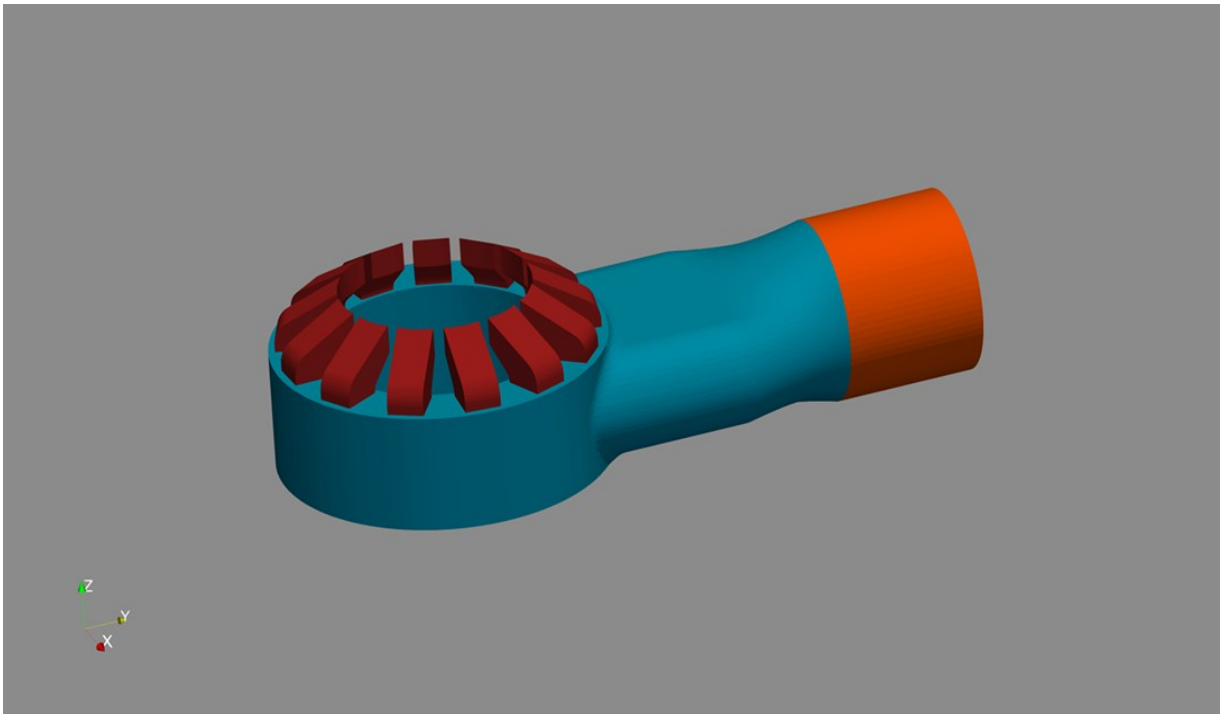

Figure 24: Opposed Piston exhaust geometry

To capture sharp features, both intake and exhaust volumes were split into sub-parts: detailed duct sections were meshed with blockMesh, while rings and runners were meshed with snappyHexMesh. Short extrusion section near outlet patch was also created with blockMesh (orange part in Figure 24). Splitting is important because snappyHexMesh may struggle at well corners, especially at intake/exhaust–cylinder interfaces where mesh quality is critical.

The intake volume comprises 15 ducts, six of which include hydrogen straws. Two additional ducts (marked 1–2 in Figure 23) have shapes that differ from the others; they were further split and meshed with blockMesh and snappyHexMesh separately.

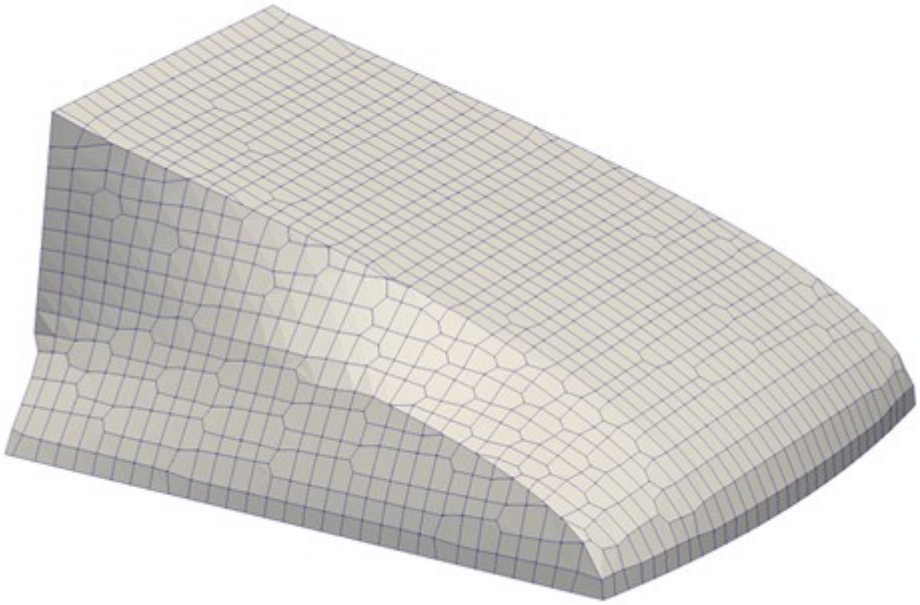

Figure 25: special duct part meshed with snappyHexMesh

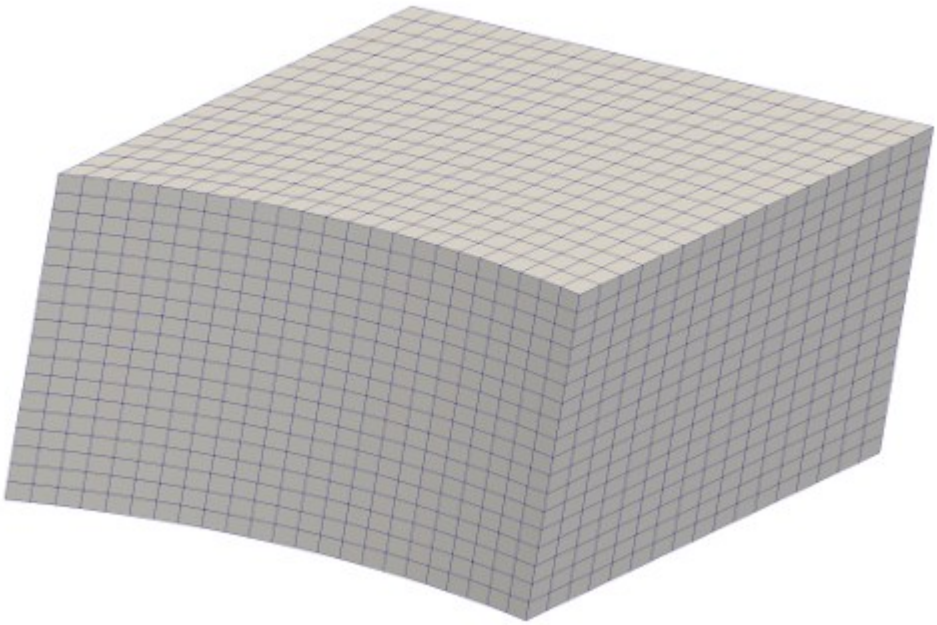

Figure 26: special duct part meshed with blockMesh

Ducts with injector features were first generated with blockMesh; injector pockets were then removed using snappyHexMesh boolean operations (Figure 27–Figure 28).

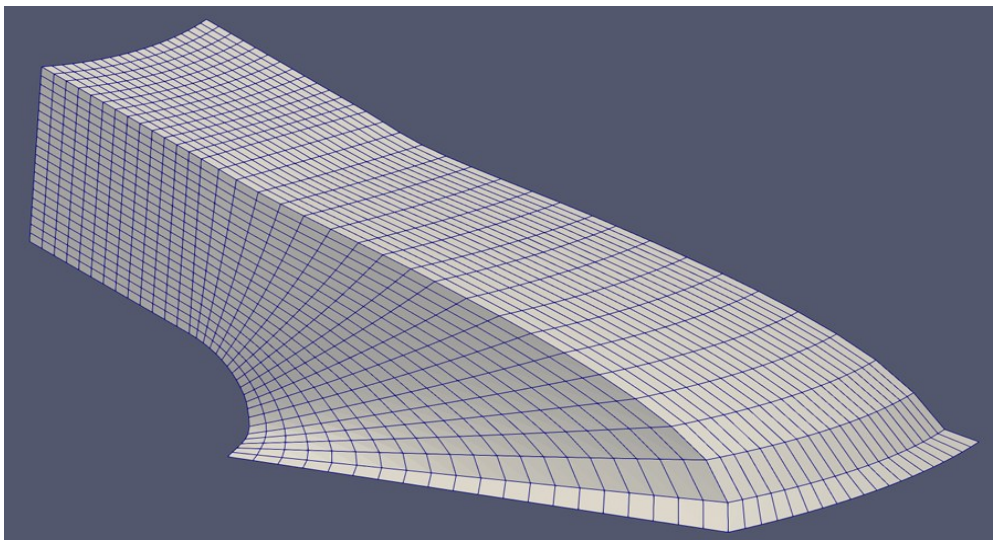

Figure 27: intake duct before injector adding

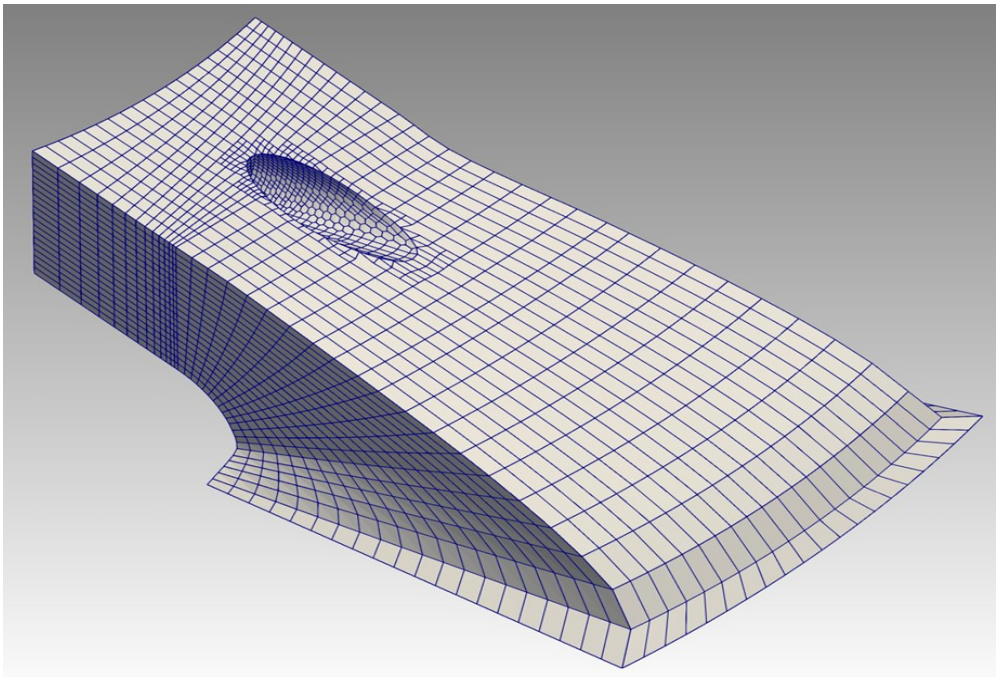

Figure 28: intake duct after injector adding

Individual meshes were combined using stitchMesh (OpenFOAM native). Because stitchMesh can degrade interface quality, it was used sparingly; in this case, the resulting maximum skewness ≈ 4.6, slightly above the OpenFOAM guideline (~4), but acceptable for the subsequent simulations.

The cylinder mesh was generated with blockMesh (Figure 29–Figure 30). LibICE utilities automated parts of the workflow by exposing a small number of parameters to reach the target refinements.

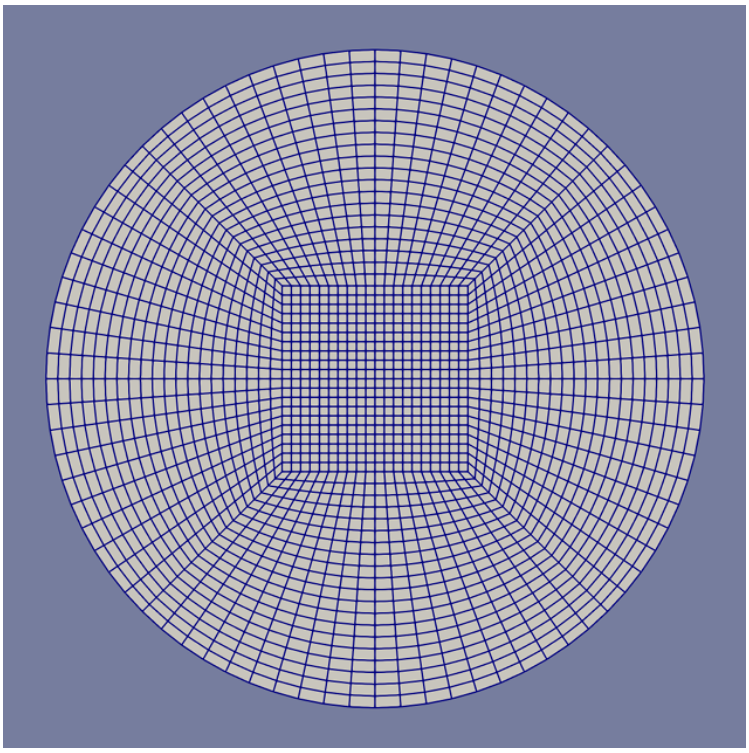

Figure 29: cylinder mesh section

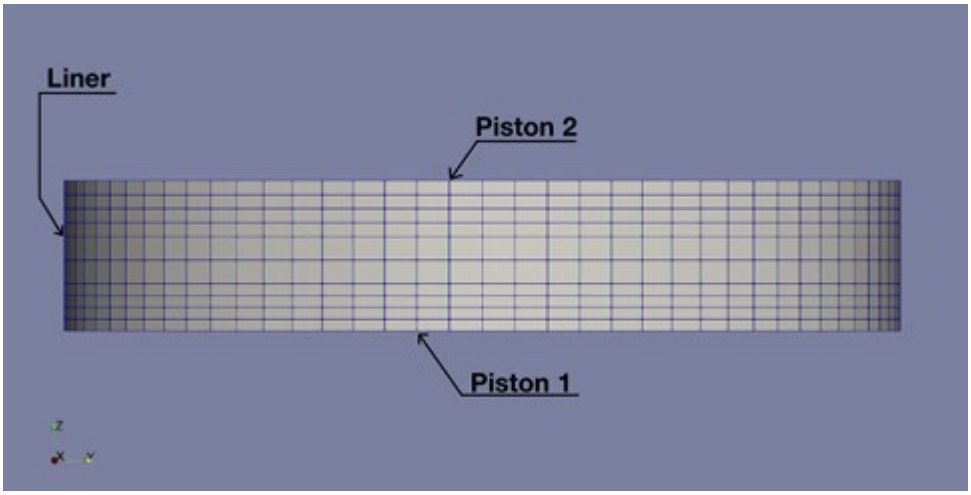

Figure 30: cylinder mesh side view

2. Cylinder mesh motion

Cylinder mesh motion was handled entirely by LibICE. A layering (add/remove) strategy was used: selected cell vertices are moved so that, for example, as the pistons travel toward BDC, cell heights increase along a given axis (here, Z). When a threshold is reached, a new layer is inserted; during contraction, layers are removed once height drops below a threshold. This approach is simple to run-time manage and minimizes pre-processing, but it only applies to structured cells—hence the layered design in Figure 30.

A second motion option is Laplacian smoothing, which moves vertices without adding/removing cells. It is used where intake/exhaust volumes interact with the cylinder and ACMI (Arbitrary Cyclic Mesh Interface) interfaces are active—i.e., from exhaust-port opening up to intake-ports closing. ACMIs require similar face areas on the coupled patches; since ACMIs use area-weighted interpolation, a large area mismatch (e.g., ratio > ~1.5) can yield non-conservative fluxes and degraded accuracy.

3. Pre-processing

3.1. Boundary conditions

Before running the simulations, boundary conditions (BNDs) from the GT-Power 1D model were applied in OpenFOAM.

OpenFOAM provides several BND types; the most relevant here are:

- fixedValue — Enforces a prescribed value on the patch. It is faithful to the provided time history but can create gradients normal to the boundary if the internal-cell value differs significantly, especially when velocity tends to zero while flow exits the domain. In that case, pressure and velocity on the patch differ from the internal field and divergence is likely, as sketched in Figure 31.
- inletOutlet — Distinguishes inflow from outflow. For outflow, the condition reverts to zeroGradient so that no spurious divergence arises; for inflow, the BND values are imposed. While this cures divergence, it may underperform physically if used indiscriminately because BNDs are enforced only for inflow.

To represent the investigated physics, total pressure and mass-flow-rate (MFR) BNDs were used. Total pressure is an inletOutlet type but still predicts flow correctly across the patch [19].

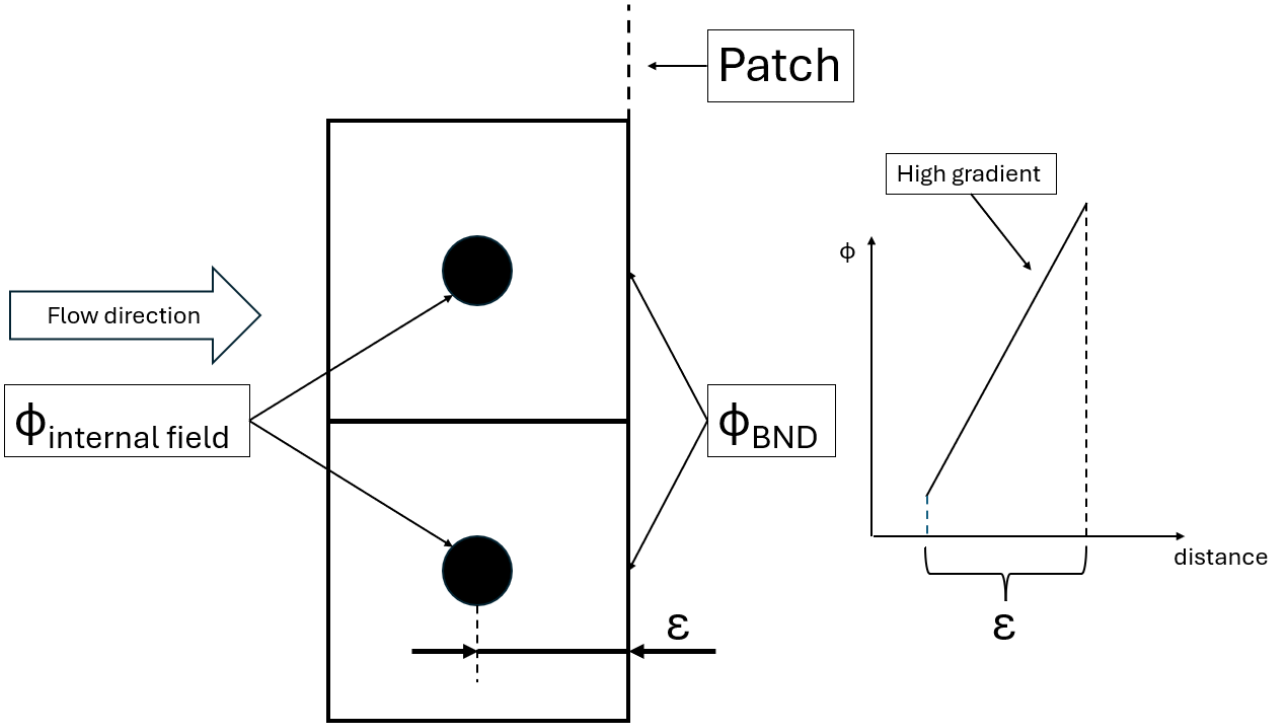

Figure 31: fixedValue gradient issue

Applied BNDs are:

- Inlet — MFR condition. Using MFR (instead of static/total pressure) eases cycle-to-cycle convergence and reduces run time to reach mass convergence in the cylinder. A crank-angle static-temperature profile was imposed
- Outlet — Total pressure condition. Total pressure is recommended so that the internal flow matches the target gas dynamics. Using static pressure here would misrepresent the exhaust behavior. A crank-angle static-temperature profile was imposed
- Walls — No-slip for velocity; fixed temperatures; turbulence wall functions as appropriate

Mixture composition**.** At the intake, only dry air enters (mass fractions 0.767 N_2, 0.233 O_2). Special care is required for injection patches. Before injection starts, the injector patch acts as a wall (no-slip, fixed temperature, turbulence wall functions). When injection is on, the patch type must switch from wall to patch. This can be coded as a custom boundary condition in C++, but that requires implementing code for every field (U, p, T, k, ω, species, etc.). The practical route is to stop the simulation at the crank angle where the change occurs, modify the BNDs, and restart. The reverse switch is applied when injection ends.

H_2 is injected via MFR at the injector inlet patches, using GT-Power time histories. 3D CFD is then used to assess how different MFR profiles affect scavenging metrics such as trapping efficiency. Figure 32–Figure 33 show two profiles that inject the same total mass; the second profile yields higher trapping.

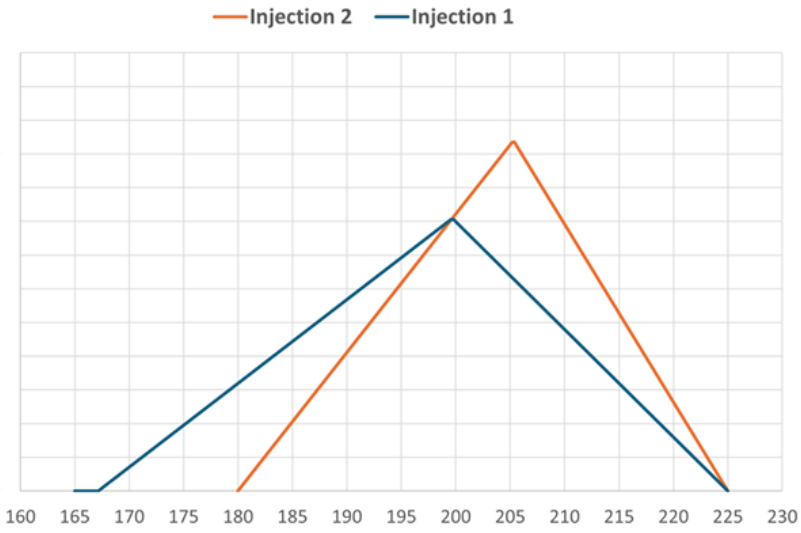

Figure 32: injection profiles

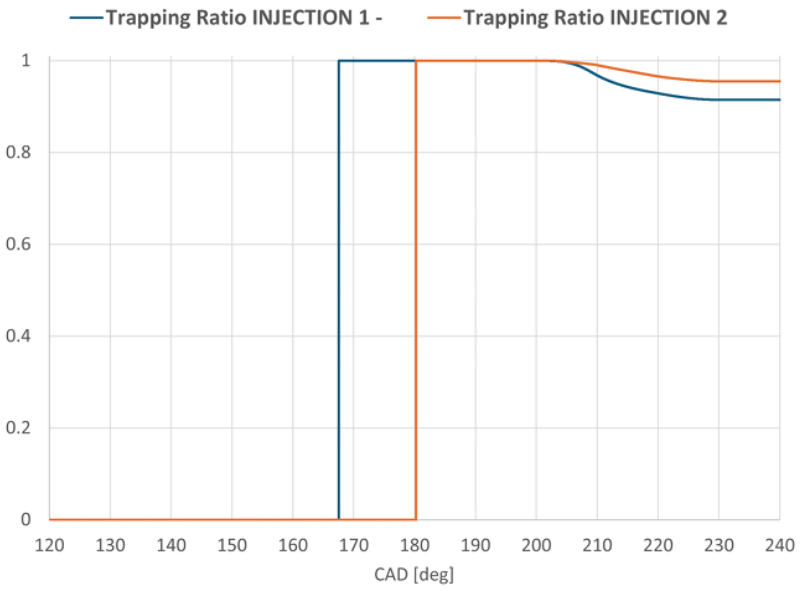

Figure 33: hydrogen cylinder trapped mass

Since OP engines lack poppet valves, minimizing fuel short-circuiting to the exhaust (and maximizing brake thermal efficiency) relies on optimizing the scavenging flow field.

3.2. Function objects

To compute performance metrics from 3D results, function objects are required. OpenFOAM ships several native post-processing functions (volume averages, patch fluxes, etc.), but each produces outputs independently and some quantities (e.g., species mass flow through ACMI patches) are not directly available.

Therefore, I implemented a custom function object (C++ class). Instead of pasting large coded snippets into controlDict dictionary for every case (Figure 34), the class exposes a single configuration block (Figure 35) so users can select what to compute at run time.

```
writeMagU    // function name
  {
    libs      ("libutilityFunctionObjects.so");
    type coded;
    codeWrite
    #{
      // write code here
    #};
  }
```

Figure 34: user-defined function in controlDict

The new function object configuration keys are **(**Figure 35**):**

- scalarFieldsToAverage — List of volScalarField objects to volume-average (e.g., p, T, k, ω).
- calculateZoneMassesAndVolumes *(bool)* — If true, computes total mass (per species) and total volume for each cellZone. Species lists come from thermophysicalProperties (entry speciesToAverage).
- zonesToAverage — Names of cellZones over which to compute the requested quantities.
- calculatePatchFluxes *(bool)* — If true, computes total mass flux and species mass flux across patches in patchesFluxList.

```
fieldsToAverage1
  {
     type  fieldsAverage_Claudio;
     libs ("libfieldsAverage_ClaudioFunctionObject.so");
     scalarFieldsToAverage ("p" "T" "k" "omega");
     calculateZoneMassesAndVolumes true;
     zonesToAverage ("cylinderZone" "intakeZone" "exhaustZone");
     calculatePatchFluxes true;
     patchesFluxList ("IntakePortOpen" "linerIntakeOpen" "ExhaustPortOpen" "linerExhaustOpen");
  }
```

Figure 35: new function object

With this function object, users can add/remove averaged fields and enable species masses or patch fluxes as needed—without duplicating code. Results are written at each time step as CSV files, ready for subsequent analysis.

3.3. New Utilities

In cold-flow simulations of two-stroke engines, key performance metrics require distinguishing the fresh charge entering from the intake from the residual gases carried over from the previous cycle. Although molecules such as N_2 are chemically identical, they must be tagged by origin (fresh vs residual) to quantify scavenging efficiency and short-circuiting.

A practical approach is to duplicate the transported species and treat each pair as fresh and residual (burnt) tracers (e.g., N_2/N_2_burnt, O_2/O_2_burnt, H_2/H_2_burnt). This labeling allows the contribution of fresh air versus residuals to be computed at any time and location (e.g., the fresh-air fraction of O_2, N_2, H_2). Then, EGR (residual-gas) fraction follows directly as the sum of the "burnt" components (typically CO_2_burnt + H_2O_burnt + N_2_burnt, normalized by the total). If backflow/reversion occurs, the model naturally tracks burnt species convected from the cylinder back into the intake.

For multi-cycle simulations, species tags must be switched when a new cycle starts: burnt species at the intake become fresh, and fresh species at the exhaust become burnt. Moreover, multi-cycle combustion is usually not run; pressure and temperature are re-initialized each cycle. To support these operations, I developed two new utilities: fieldOperation and fieldsMultiplier.

fieldOperation has three mandatory and one optional input:

- volScalarFields: list of volScalarField objects to manipulate.
- operation: operation to perform — sum, subtract, multiply, or divide.
- targetName: name of the resulting object; can overwrite one of the two inputs.
- zoneNames *(optional)*: list of cell-zone names on which to apply the operation (otherwise the whole domain).

For each computational cell, the utility applies the selected operation to the two input fields and writes either a new field or overwrites the chosen input (targetName). The first input is treated as master and the second as slave; this distinction is mandatory for subtract and divide.

FieldsMultiplier utility multiplies each cell value of a given scalar field by a coefficient equal to the ratio between a target volume-averaged value and the current volume-averaged value. Target and current averages are read from a dictionary (see Figure 36). Entries include the list of volScalarField objects and, optionally, zoneNames. If zoneNames is provided, the dictionary's regions section is used; otherwise, the global section applies.

The simulation must be stopped to run these utilities. A planned improvement is to enable their use at run time.

```
// SUBDICTIONARY "regions" USED WHEN “-zoneNames" OPTTION IS USED
    regions
    {
     cylinderZone   // ZONE NAME (KEYWORD)
     {
      volAverageValues   // KEYWORD
      {
        p 4.94907e5;    // volScalarFields volAveraged target values
        T 1480.9736307517;
      }
     }
     zone2
     {
          volAverageValues
          {
                    H2O 0.01;
                    CO2 0.19;
          }
     }
    }

    // IF "-zoneNames" IS NOT SPECIFIED THEN "global" SUBDICTIONARY IS USED
    global
    {
     ft     0.054;   // volScalarFields volAveraged target values
     egr    0.165;
    }
```

Figure 36: fieldsMultiplierDict dictionary

4. Numerical discretization

One of the most critical aspects of 3D CFD is the discretization of scalars. In OpenFOAM, this is defined in the fvSchemes dictionary, where the user selects the schemes for each term. The most impactful choice is typically the convective discretization of the transport equations (notably the velocity field). To preserve accuracy, higher-order schemes are preferred.

OpenFOAM offers many options, but three are commonly used: limitedLinear, linearUpwind, and upwind:

- limitedLinear and linearUpwind are second-order schemes
- upwind is first-order

In this work, second-order schemes were adopted for all fields except turbulence, for which first-order upwind was chosen. At this stage, no wall layers were included during meshing; to prevent spurious turbulence oscillations (which could trigger divergence), the upwind scheme was used for turbulence quantities.

5. Cold flow results (scavenging)

The cold-flow case was run without fuel injection to characterize scavenging independently of injection effects. Because injection can alter the intake flow and its distribution, excluding it yields results more representative of the geometry-driven behavior. The simulation started before Exhaust Port Opening (EPO). A swirl ratio of 1.44 was imposed as the initial velocity field, and the standard k–ε turbulence model was used.

A primary check is the cylinder total mass, which should remain close to the initial mass specified by the boundary conditions. Figure 37 shows the cylinder-mass trace; mass conservation is satisfied, with a difference of 1.13%, indicating that the 1D BNDs capture the intended gas dynamics.

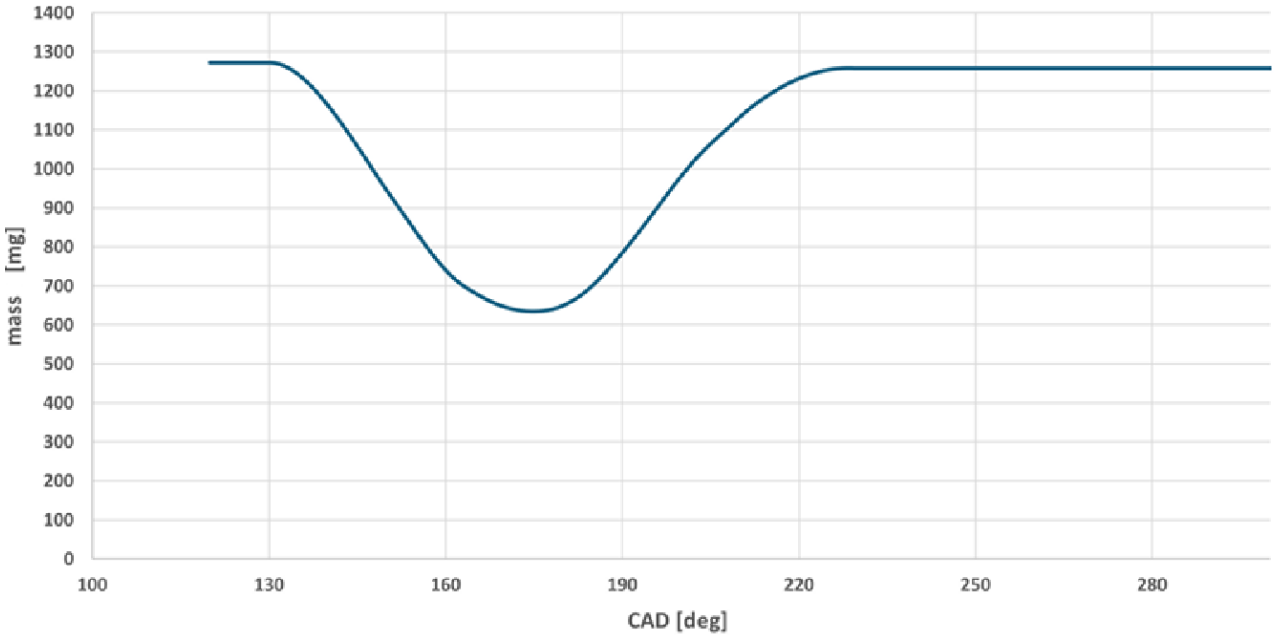

Figure 37: cylinder total mass

Scavenging performance is characterized by the following metrics:

- Delivery Ratio (DR): fresh mass delivered to the cylinder divided by a reference mass (air density times displacement).
- Trapping Efficiency (TE): ratio of fresh mass trapped in the cylinder to the total fresh mass delivered. It quantifies how much of the delivered mass is retained; values < 1 indicate short-circuiting to the exhaust.
- Charging Efficiency (ChE): ratio of trapped fresh mass to the reference mass
- Scavenging Efficiency (SE): ratio of trapped fresh mass to the total cylinder mass

Figure 38 plots these parameters versus Delivery Ratio. The results are:

- TE = 0.961, meaning ≈4% of the delivered fresh mass is lost to the exhaust.
- DR = 1.0385 (slightly above unity).
- ChE = 0.998
- SE = 0.908, indicating that at end of scavenging the residuals from the previous cycle are 9.2% of the cylinder mass.

Up to DR ≈ 0.78, the process is close to the perfect displacement mixing model (one of the limiting scavenging models).

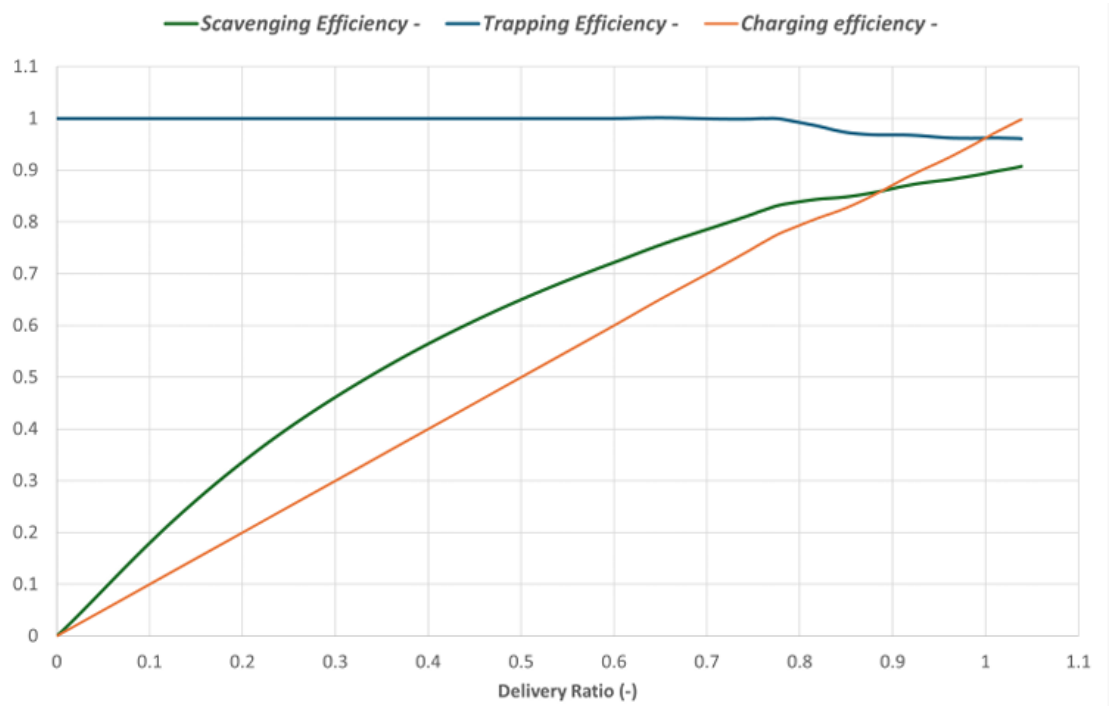

Figure 38: scavenging performance parameters

Figure 39 shows the residual-gas mass at the exhaust versus the residual-gas mass in the cylinder. Up to a cylinder residual fraction of 0.26, the behavior follows the perfect mixing model and no short-circuit is observed.

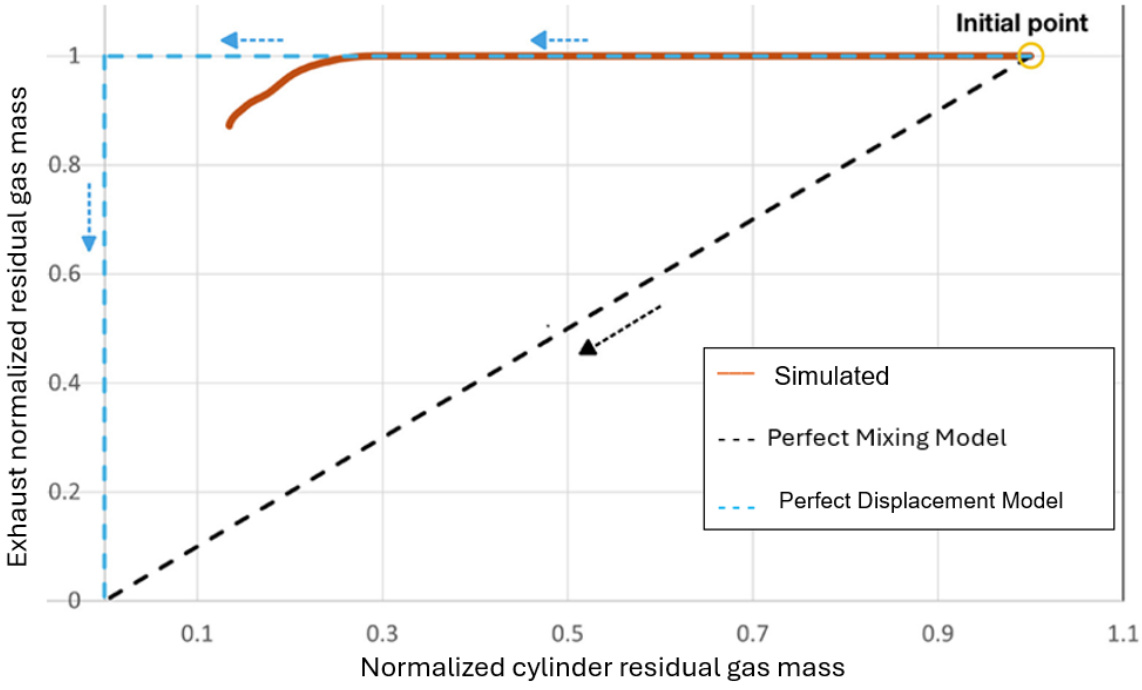

Figure 39: scavenging plot

By comparing the 3D CFD scavenging results with the 1D model (Figure 40), both approaches predict similar performance metrics—Delivery Ratio, Charging Efficiency, Trapping Efficiency, and Burned-Gas Fraction—indicating good agreement between the two models.

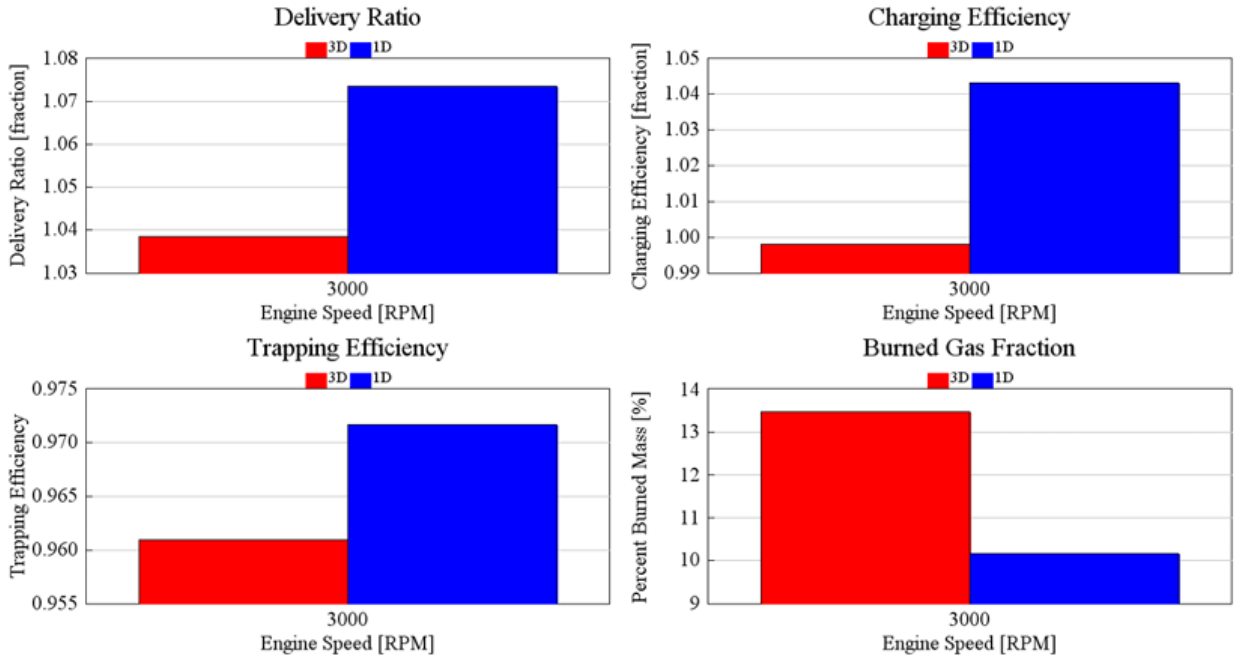

Figure 40: 1D-3D scavenging performance comparison

6. Combustion results

As noted above, hydrogen is introduced by imposing an MFR profile at the injector inlet patches. Analysis of the 3D results shows that the in-cylinder fuel is well mixed before combustion begins. Figure 41 reports the λ (lambda) distribution on a plane normal to the Z-axis through the cylinder mid-height. This outcome, consistent with previous works, supports the approximation of numerically imposing a uniform fuel distribution.

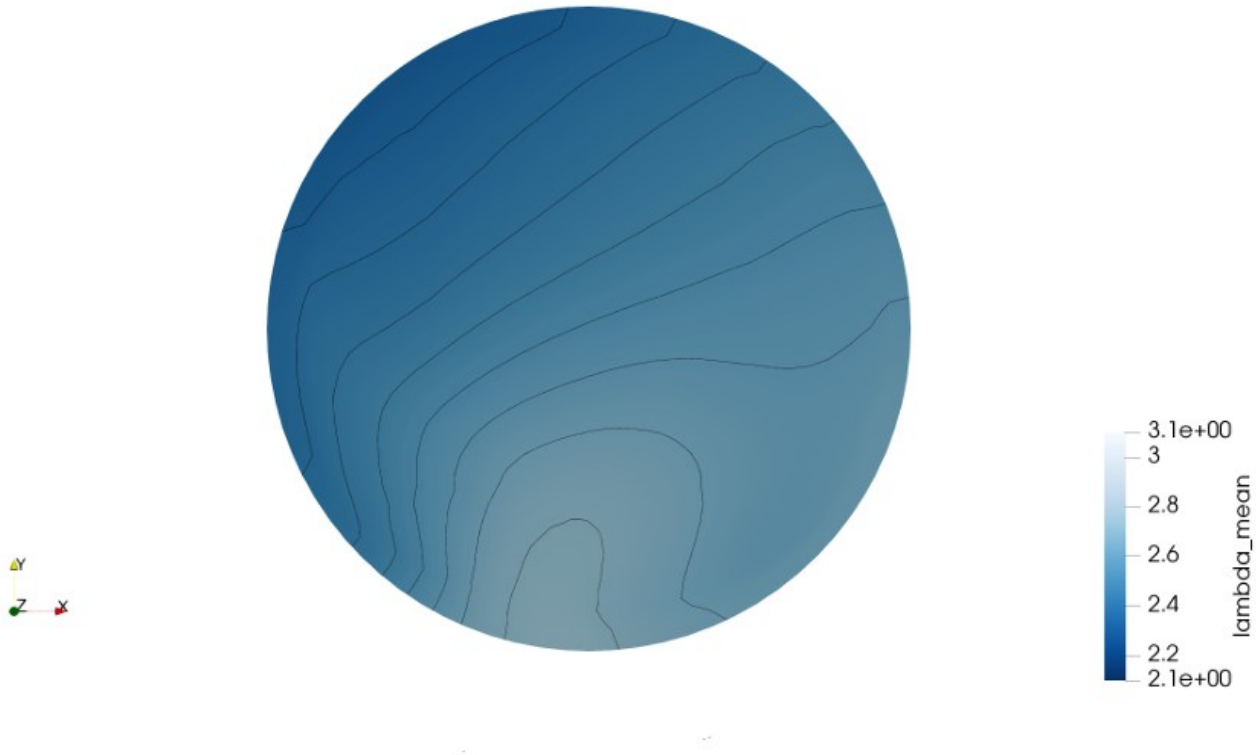

Figure 41: lambda distribution at cylinder midline at 330 CAD

Within the combustion analysis, this assumption reduces computational cost: explicitly resolving injection requires very small time steps to limit the Courant number, which increases runtime. Because the velocity field during the burn phase is driven mainly by pressure gradients (BCs and geometry), multiple combustion simulations at different lambda values can be performed simply by changing the fuel mass fraction in the cylinder. This capability is natively supported by LibICE, whose combustion model allows prescribing a target fuel mass fraction. The uniform-mixture assumption is particularly valid for the very lean mixtures typical of H_2-fueled ICEs.

For the combustion runs, only the cylinder was simulated. Because layering is used for mesh motion, the only way to increase resolution is to globally refine the cylinder mesh; local refinements are not allowed, otherwise the mesh would lose its structured layout (see Figure 42).

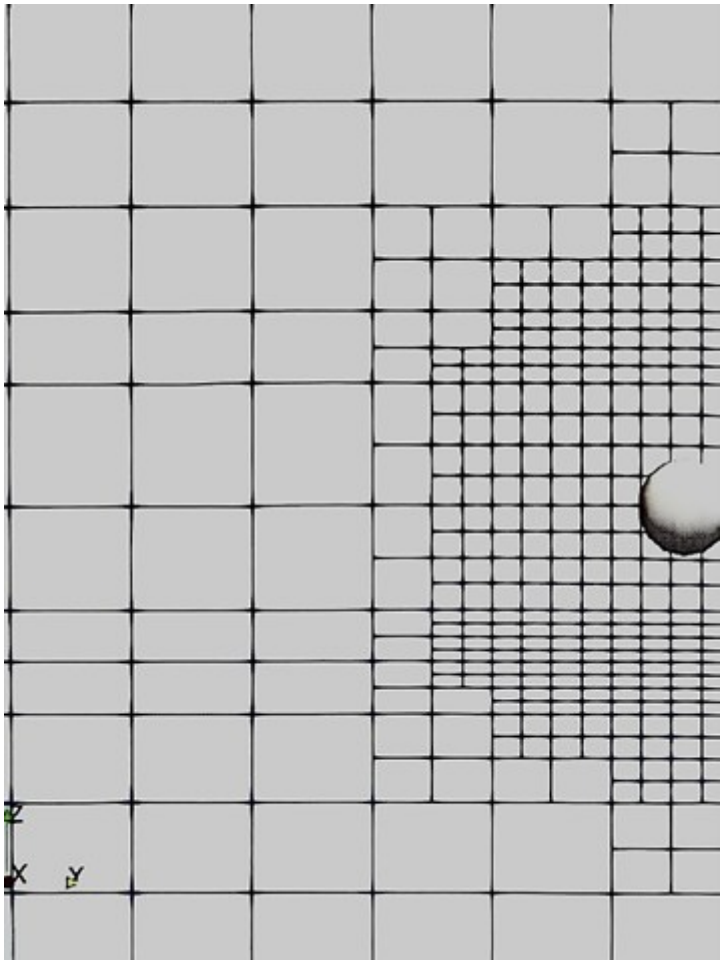

Figure 42: spark refinement layering issue

A global λ = 2.4 (equivalence ratio φ = 0.416) was imposed. To promote a more complete burn, two spark plugs were positioned on opposite sides of the chamber. The ignition model selects cells within spherical kernels where the regress variable will be consumed. The Weller combustion model employed here evolves the regress variable (rather than a progress variable); the two are complementary:

$$Progress\ Variable = 1 - Regress\ Variable$$

Three spark-advance (SA) values were tested: 20, 30, and 35 CAD bTDC. Figure 43–Figure 46 report IMEP, MFB10–90, adiabatic efficiency, and combustion efficiency. As SA increases, IMEP rises and turbulent burn duration (MFB10–90) decreases, because earlier ignition raises cylinder pressure and temperature at comparable CAD, accelerating both laminar and turbulent combustion. While higher pressure tends to reduce laminar flame

speed, the temperature effect dominates, so higher SA generally improves combustion performance.

At the same time, SA cannot be increased arbitrarily: excessive pressure before TDC raises piston work losses (gas forces decelerate the piston), which can offset gains. This trade-off is evident in Figure 43, where SA = 35 CAD does not increase IMEP relative to 30 CAD, although combustion efficiency is higher for 35 CAD (Figure 46).

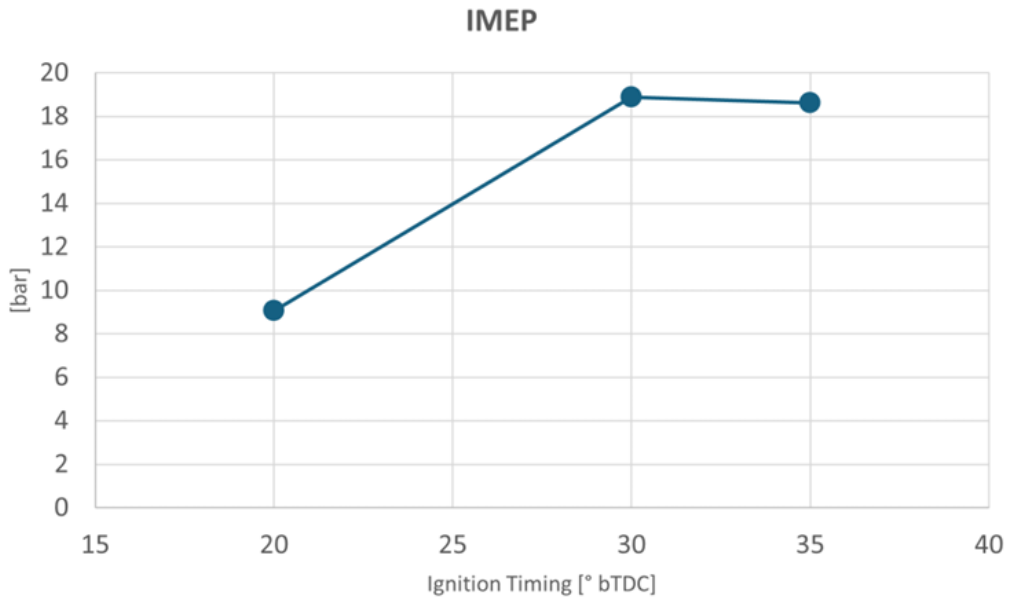

Figure 43: Indicated Mean Effective Pressure (IMEP)

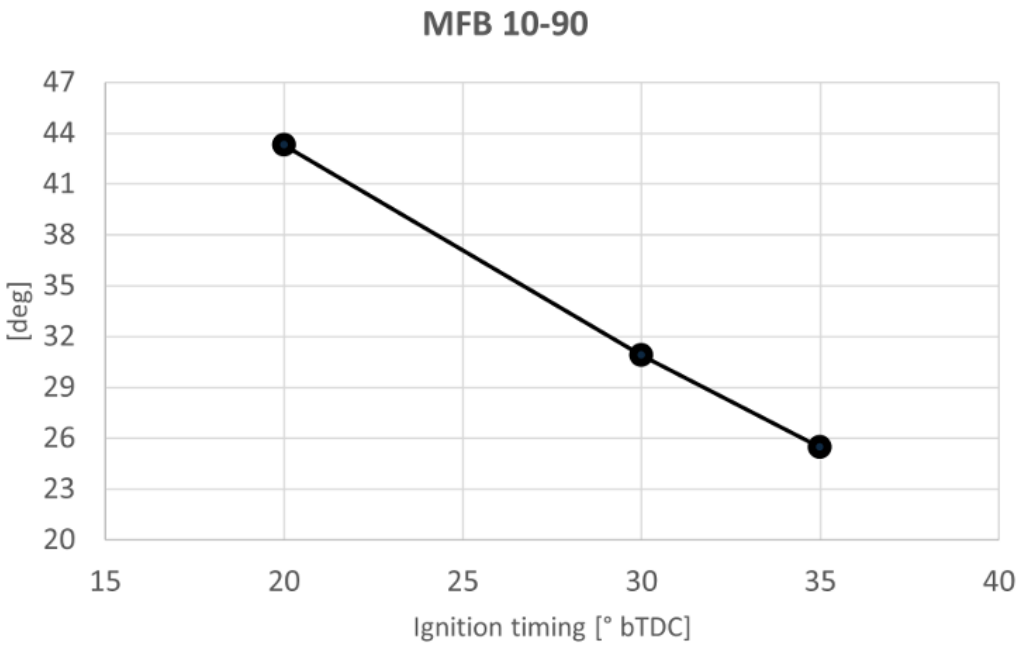

Figure 44: 10-90% Mass Fuel Burnt Duration

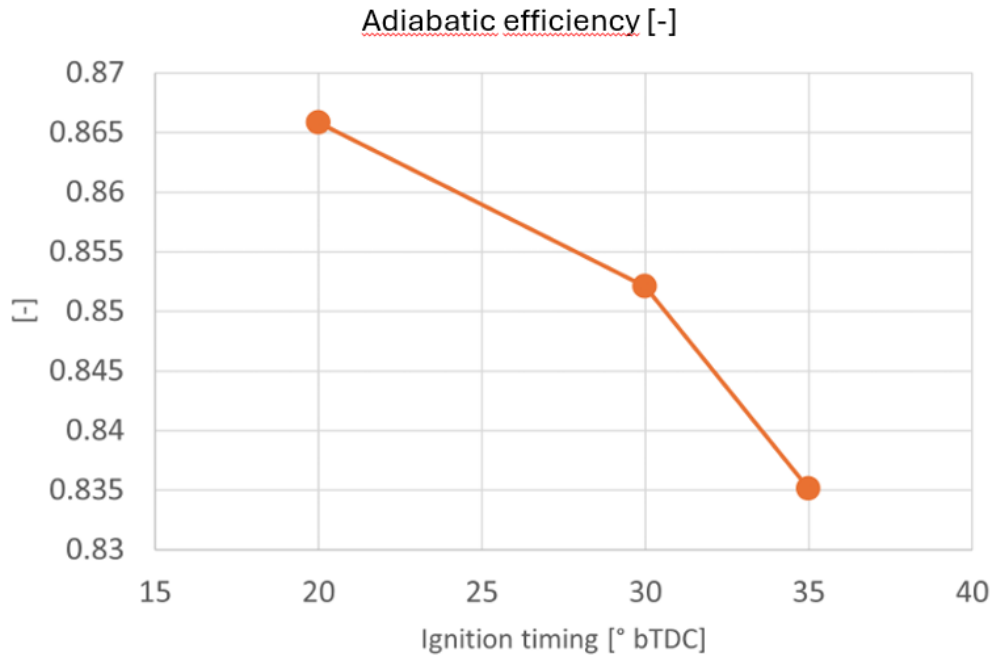

Figure 45: adiabatic efficiency

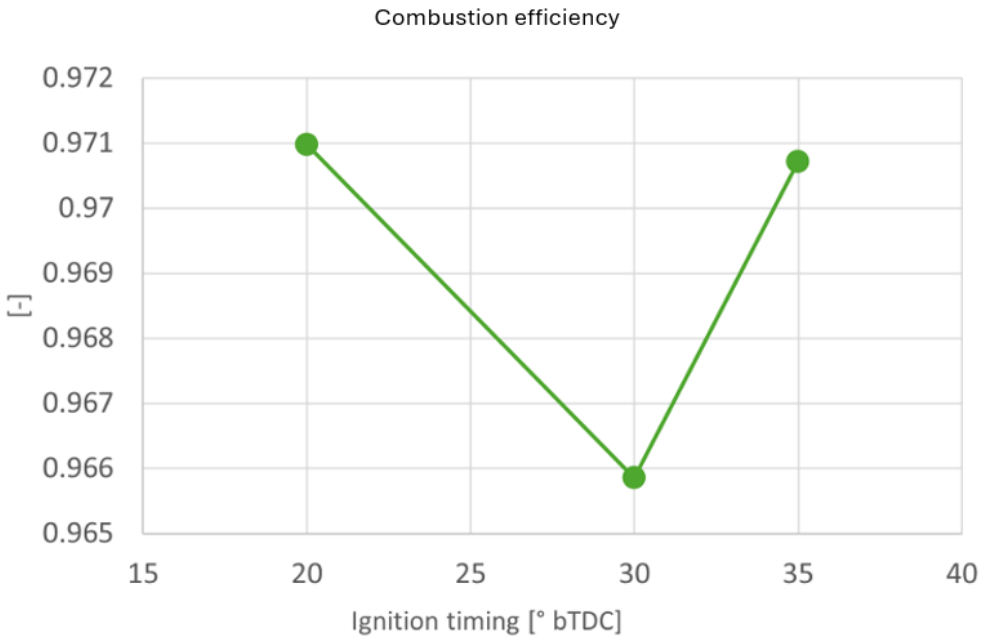

Figure 46: combustion efficiency

As shown in Figure 47, the combustion model correctly reflects thermodynamic-efficiency trends. Comparing SA = 30 and SA = 35 CAD bTDC, the latter yields higher peak temperatures while reaching similar end-of-burn temperatures to the 30 CAD case; consequently, SA = 35 features higher thermodynamic efficiency. Conversely, SA = 20 CAD exhibits lower peak temperatures and a higher late-cycle temperature, which together lead to lower thermodynamic efficiency.

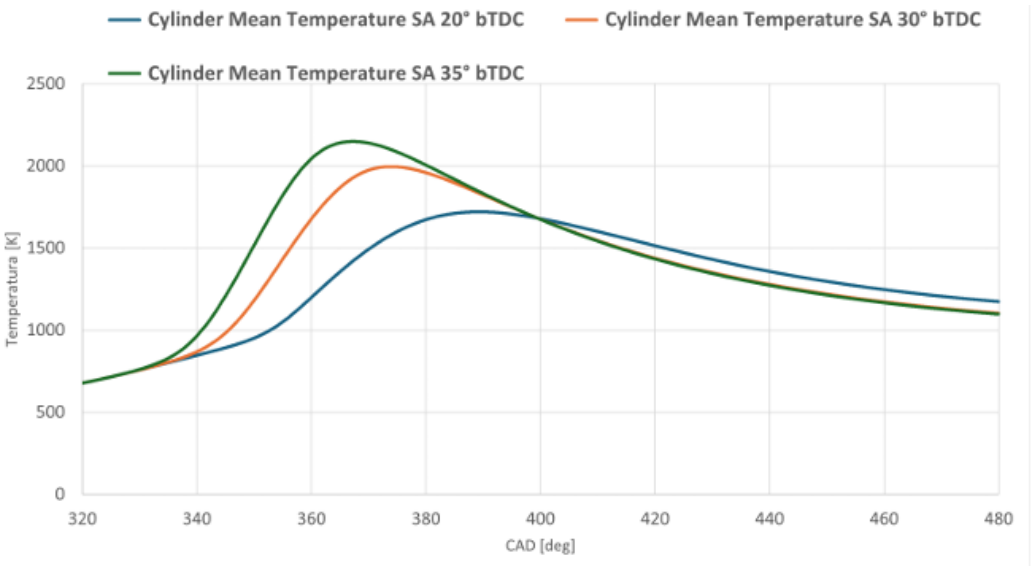

Figure 47: cylinder mean temperature

NO_x formation is strongly linked to temperature: the higher the temperature, the greater the NO_x production. The model captures this behavior using the Zeldovich mechanism, as illustrated in Figure 48.

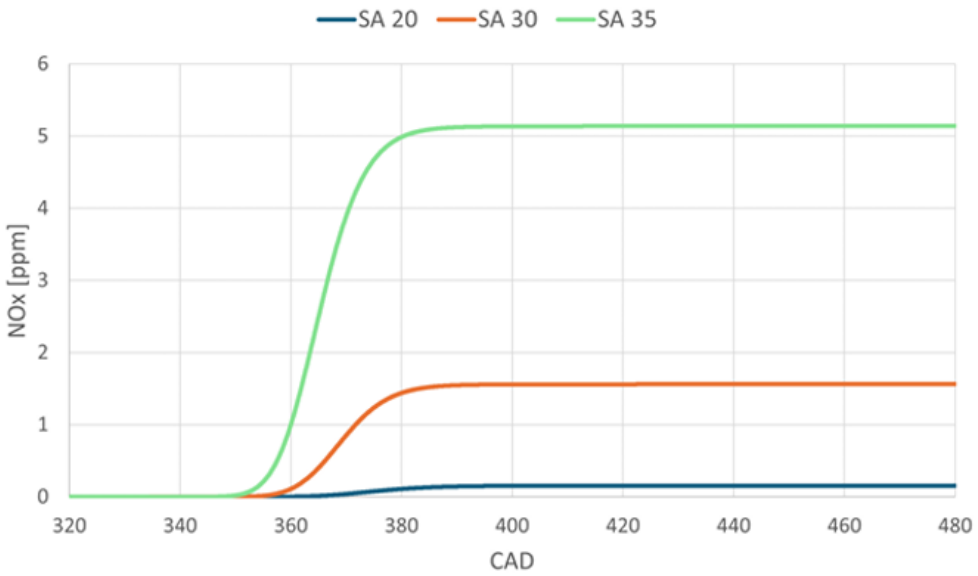

Figure 48: NOx production

7. Conclusions

This work presented a virtual investigation of a hydrogen-fueled opposed-piston spark-ignition ICE. A 1D GT-Power model provided boundary conditions, and 3D CFD simulations were conducted in OpenFOAM 8.

Static meshes were generated with blockMesh and snappyHexMesh, ensuring adequate geometric fidelity and overall quality while placing refinement where needed. LibICE provided the structured cylinder mesh and layering mesh-motion capability.

A new function object was implemented to compute volume averages and mass fluxes (including species) through patches and to write results as CSV. A fresh/burnt species labeling concept was adopted to distinguish residuals from fresh charge across cycles; two utilities were developed to automate species re-tagging and field operations for multi-cycle pre/post-processing.

The cold-flow (scavenging) simulation showed good agreement with the perfect-displacement model up to Delivery Ratio ≈ 0.78. Two injection phases were examined to assess trapping efficiency; profiles with the same injected mass produced similar trends, with one profile yielding slightly higher trapping.

Combustion simulations at $\lambda = 2.4$ (uniformly imposed, a valid assumption for very lean H_2 mixtures) were run with SA = 20, 30, 35 CAD bTDC. Results indicate that IMEP is maximized at SA = 30 CAD, while SA = 35 CAD provides higher thermodynamic/combustion efficiency but no further IMEP benefit, and turbulent burn duration (MFB10–90) remains comparable to values typical of gasoline engines under similar conditions.

CHAPTER 4

Spray 3D CFD model calibration for DISI ICEs applications

During my PhD abroad, I worked on different fuels for internal-combustion engines. The objective was to investigate—via 3D CFD—the influence of different fuels on ICE behavior. Experimentally calibrated spray models were therefore required, which was the focus of this activity and carried out in STAR-CCM+. No sensitive data can be shared.

The available experimental inputs were:

- droplet-size probability density function
- post-processed images of the injected fuel in a spray-bomb box
- mass-flow-rate profiles
- thermodynamic conditions in the spray bomb (temperature, pressure)
- Sauter Mean Diameter (SMD)
- injector-hole diameters
- liquid-column diameter
- spray-cone opening angles and axis directions
- definition of the fuel surrogate

Injection was modeled with a Lagrangian approach, in which the liquid phase (injection) is separated from the Eulerian gas phase. Injected particles are computed as parcels, i.e., computational objects that represent groups of droplets sharing properties such as size and temperature. This avoids resolving the internal nozzle flow and reduces computational cost; however, initial droplet properties (size, temperature, velocity) must be prescribed and are subsequently evolved by primary/secondary breakup models.

1. Meshing

Mesh generation used STAR-CCM+'s automated mesher: given the STL surface, the software creates a discretized volume. To remain consistent with other project requirements, cell type (polyhedral) and target size were chosen to meet a specific resolution.

2. Model set up

STAR-CCM+ provides a structured workflow to configure the spray simulations. The key models enabled were:

- multi-component gas phase
- multi-component Lagrangian phase
- Eulerian–Lagrangian coupling via source terms (drag, turbulence, heat transfer)
- k–ε turbulence model
- KH-RT secondary-breakup model

3. Initial droplets properties

As liquid flows through the injector cavities, several phenomena shape the spray. Droplets detach from the external liquid column near the orifices—this is the primary breakup.

Based on droplet size and the axial distance at which detachment occurs, four primary-breakup regimes are identified (Figure 49):

- Rayleigh: low flow velocity; droplet diameters exceed the liquid-column diameter; breakup occurs relatively far from the nozzle
- Laminar/turbulent (first wind-induced): higher velocities than Rayleigh; droplet diameters comparable to the liquid-column diameter
- Second wind-induced: droplet diameters smaller than the liquid-column diameter; breakup occurs much closer to the exit
- Atomization: very high injection pressure and nozzle-exit velocity; droplets are much smaller than the liquid column and breakup is nearly instantaneous at the exit

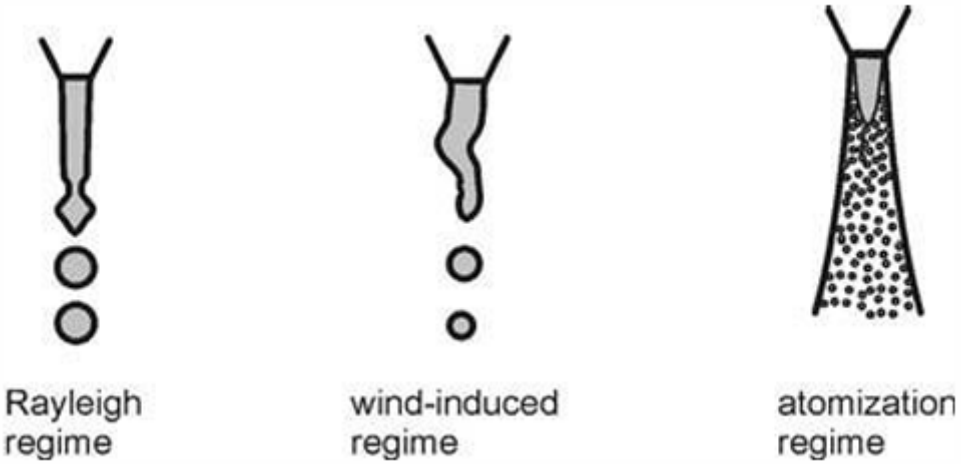

Figure 49: primary breakup regimes [20]

A common nondimensional indicator is the liquid Weber number,

$$We_l = \frac{u^2 D \rho_l}{\sigma}$$

where u is the relative particle–gas speed, D the droplet diameter, ρ_l the liquid density, and σ the surface tension. Higher injection velocity increases We_l and thus atomization. An analogous gas Weber number can be defined by substituting gas properties, highlighting that breakup depends on both liquid and gas states. For DISI applications, with injection pressures on the order of 300 bar [21], the relevant regime here is atomization.

Using a Lagrangian approach, the user must prescribe the primary-breakup droplet size leaving the liquid column. In this work, two prescriptions were employed: the Blob method and the Rosin–Rammler distribution.

The Blob method imposes a single droplet size at release (Figure 50). This reflects the experimental difficulty of measuring droplet characteristics directly at the nozzle exit due to the dense liquid core.

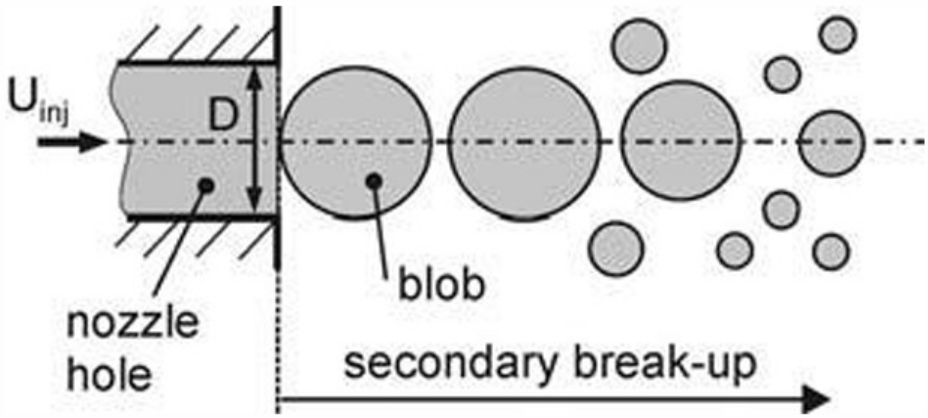

Figure 50: Blob Primary Breakup [20]

The Rosin–Rammler method injects droplets with diameters sampled from a cumulative distribution function (Equation 4), where D_{ref} is the reference diameter, n the shape parameter, and D the droplet diameter:

$$F(D) = 1 - \exp\left[-\left(\frac{D}{D_{ref}}\right)^{n}\right]$$

Equation 4: Rosin-Rammler Cumulative Distribution Function

For both primary-breakup options, spray-cone geometries were defined and droplets were released inside the cones; the opening angle was set to the experimental value.

Initial droplet velocities were estimated from Equation 5,

$$\dot{m} = \rho_l A_l |v|$$

Equation 5: injection mass flow rate law

where ρ_l is the liquid density, A_l the cross-sectional area of the liquid column, and |v| the velocity magnitude normal to A_l. This is a simplified relation that neglects internal-nozzle perturbations and cavitation, but it is reasonable here because experimental $\dot{m}$ and nozzle-section area were available. The fuel temperature was imposed directly according to the operating point.

4. Fuel surrogate model

Real fuels contain hundreds of components. In CFD, two main strategies are used: single-component and multi-component surrogates.

Both create a virtual (surrogate) fuel. In the single-component approach, one pseudo-component is defined whose properties are averaged from the real fuel. In the multi-component approach, the surrogate includes the most relevant species of the real fuel (e.g., iso-octane, ethanol), with mass fractions chosen to match bulk properties of the target fuel (thermal conductivity, density, saturation pressure, etc.).

5. Calibration methods

This section summarizes best practices identified during the activity.

CFD guides recommend placing injection source points at least two cells away from walls. Initially, I positioned them as close as possible to the injector coordinate origin, but the results showed non-physical behavior—likely due to wall interactions. After moving the sources farther from the walls, the issue disappeared.

Measured droplet PDFs are typically reported as:

- volume-based: fraction of particle volume (over total volume) per diameter bin
- number-based: fraction of particle count (over total count) per bin
- mass-based: fraction of particle mass (over total mass) per bin

This distinction matters because it affects the Sauter Mean Diameter evaluation. In our dataset, the only PDF metadata was the axial distance from the nozzle where it was measured; no other details were given. The SMD computed with Equation 6 matched the datasheet value, indicating that the provided PDF was number-based:

$$SMD = \frac{\int_0^\infty D^3 * PDF(D)dD}{\int_0^\infty D^2 * PDF(D)dD}$$

Equation 6: Sauter Mean Diameter number based

5.1. Turbulence initial conditions

The initial Eulerian turbulent kinetic energy (ITKE) and its dissipation rate (ITKED) are critical. Higher TKE amplifies surface instabilities and promotes droplet breakup. This is consistent with the KHRT model, which accounts for Lagrangian–Eulerian coupling via source terms. ITKED also influences breakup by controlling how turbulence decays over time.

The default ITKE and ITKED values in STAR-CCM+ did not reproduce experimental results. Moreover, they should not be treated as calibration parameters: since the Eulerian phase is intended to start from a quiescent state, the initial values should be set as low as practicable.

In practice, the TKE/TKED ratio was the key parameter. With TKED fixed to unity, setting ITKE ≈ 1 provided good agreement with the experimental evidence.

5.2. Primary Breakup: Blob vs Rosin-Rammler

The first simulations applied the experimental PDF directly as a Rosin–Rammler distribution. This is reasonable only when the chamber back-pressure is close to ambient (~1 bar). At higher pressures, the predicted spray became non-physical: the jet tended to collapse, with suppressed expansion and evaporation—consistent with observations in [22].

To obtain better correlations, a larger initial droplet size must be used. However, the injected droplets interact with the gas, which induces instabilities and subsequently breaks them into two or more child droplets—this is referred to as secondary breakup.

It is important to note that the initial size set by the primary-breakup model and the secondary-breakup intensity are not independent. As reported in [22], good agreement can be obtained either by imposing larger initial droplets with strong secondary breakup, or by using the opposite combination.

Here, several initial sizes were tested with both Blob and Rosin–Rammler prescriptions. For Rosin–Rammler, parameters were set so that the mean diameter matched the Blob size. No meaningful differences were found between the two; therefore, Blob was adopted for the remaining operating points because it is simpler to use.

5.3. Secondary breakup

The Kelvin–Helmholtz Rayleigh–Taylor (KHRT) model is widely used at high Weber numbers. It combines two sub-models [23]:

- Kelvin–Helmholtz: instability driven by slip velocity between gas and droplet, shearing ligaments off the parent
- Rayleigh–Taylor: instability driven by droplet acceleration, which can shatter the droplet

Breakup occurs when the acting instabilities persist longer than each sub-model's characteristic time.

It was observed that KHRT calibration depends on the fuel surrogate. With a single-component surrogate and a fixed primary-breakup setup, the STAR-CCM+ default KHRT coefficients [23] matched experiments well. With multi-component surrogates, coefficient adjustment was required. A likely reason is how STAR-CCM+ treats parcels: average properties are computed from parcel mass fractions. For single-component fuel there is no averaging; for multi-component fuel, individual species evaporate at different rates, so parcel-averaged properties evolve in space and time. Under identical primary-breakup settings, the multi-component cases showed different penetration (liquid and vapor) and different evaporation than experiments. Following [24], the C_3 coefficient [23] was changed to recover the experimental trends.

6. Conclusions

The calibration of a 3D spray model for DISI applications was carried out in STAR-CCM+ using a Lagrangian framework, with parcels representing droplet groups and primary/secondary breakup governing their evolution. The model setup combined a multi-component gas phase, a multi-component Lagrangian phase, Eulerian–Lagrangian coupling through source terms, k–ε turbulence, and KHRT secondary breakup. Meshing relied on the automated mesher, and practical guidance from the CFD workflow proved essential—most notably, placing injector source points at least two cells away from walls to avoid non-physical spray behavior. The experimental dataset provided the droplet-size PDF, mass-flow histories, thermodynamic conditions, SMD, nozzle and liquid-column diameters, spray-cone geometry, and the fuel surrogate definition; consistency checks confirmed that the supplied PDF was number-based, as the SMD computed from it matched the datasheet value.

Primary breakup was prescribed either by the Blob method and Rosin–Rammler distribution. Under chamber conditions near ambient pressure the two options were equivalent, but at elevated back-pressure the direct use of the experimental Rosin–Rammler PDF led to non-physical collapse of the jet. More reliable agreement was obtained by increasing the initial droplet size, acknowledging that the initial size and the intensity of secondary breakup are not independent: gas–liquid interaction at release triggers instabilities that fragment the parent droplets into children. After testing several initial sizes, Blob and Rosin–Rammler settings, that yielded the same mean diameter, produced no meaningful differences, and the Blob prescription was retained for the remaining operating points because it is simpler to apply.

Secondary breakup required limited tuning. With a single-component surrogate and a fixed primary-breakup setup, the default KHRT coefficients reproduced the measurements. With multi-component surrogates, however, coefficient adjustment was necessary. This behavior is consistent with the parcel-averaging treatment in STAR-CCM+: when multiple species evaporate at different rates, parcel-averaged properties evolve in time and space, which in the present cases resulted in different liquid/vapor penetration and different evaporation than observed. Changing the KHRT C_3 coefficient restored the experimental trends. Overall, the study delivered a calibrated and transferable spray setup for DISI conditions in the atomization regime, clarified the role of initial droplet sizing versus secondary breakup, and highlighted the impact of fuel-surrogate choice on KHRT tuning.

The initial conditions of the Eulerian carrier phase were found to be consequential for spray development. In particular, the default STAR-CCM+ values of initial turbulent kinetic energy (ITKE) and its dissipation (ITKED) did not reproduce the measurements and should not be used as “hidden” calibration knobs. Since the gas phase is intended to start from a quiescent state, the initial turbulence level must be set as low as practicable. In practice, the controlling quantity was the k/ε time scale: fixing ITKED to unity and setting ITKE ≈ 1 yielded consistent agreement with the experimental trends. This choice stabilizes the KH–RT response (surface instabilities and secondary breakup) without artificially energizing the flow, and it should be treated as part of the physical initialization rather than as a tunable parameter.

Importantly, the calibrated setup reproduced the experimental data over different operating conditions without revising the previously established calibration coefficients, indicating a stable and reusable configuration.

Conclusions

This thesis brought together four strands of work that, taken as a whole, demonstrate a rigorous and transferable workflow for 3D CFD studies of advanced fuels, unconventional architectures, and spray modeling. Across all chapters, the common thread was to start from sound boundary conditions, select models whose assumptions match the physics, and then calibrate the smallest possible set of parameters against experiments—so that the resulting setups are predictive yet economical to run.

In the single-cylinder motorcycle study, E85 was assessed against a gasoline baseline using a premixed approach validated by 1D/3D correlation. A sweep of equivalence ratio and spark advance showed that E85 benefits from leaner operation ($\varphi \approx 0.9–0.95$) and slightly earlier phasing: combustion efficiency improved (up to ~22% vs E5 at lean mixtures) while performance penalties remained modest once calibration was retuned. The faster burn (shorter 10–90% MFB duration) stems from ethanol's higher laminar flame speed and its cooling effect, which concurrently mitigates knock risk. These results indicate E85 can replace conventional gasoline with an efficiency upside when the calibration is adapted to its properties.

Extending the analysis to the four-cylinder PFI engine with detailed chemistry confirmed—and refined—those trends. Using a skeletal mechanism within FIRE M, two levers were sufficient to match experiments: an energy-addition factor (EF) and a turbulence–chemistry interaction multiplier (TCI). E85 correlated well with EF≈38, while E5 required a higher TCI at the same EF—consistent with different combustion regimes for the two fuels. In that framework E85 delivered higher IMEP, cleaner emissions (lower CO and HC, higher CO_2, i.e., more complete oxidation), …and a reduced MFB10–90 burn duration relative to E5, at the cost of greater injected mass driven by stoichiometry—again emphasizing that fuel-dependent adjustment of a few parameters suffices, without the need for large-scale re-tuning or model changes.

The hydrogen opposed-piston study demonstrated how open-source tools can be extended to non-standard architectures while retaining industrially useful discipline. A carefully staged workflow—dedicated meshing (blockMesh/snappyHexMesh), structured cylinder motion via layering, robust boundary conditions, custom function objects for volume/flux accounting, and "fresh/burnt" tracer utilities—yielded scavenging predictions that track limiting-model behavior and align with 1D trends. Combustion runs at very lean mixtures ($\lambda \approx 2.4$) highlighted the phasing trade-off typical of H_2 SI: advancing the spark improves burn speed and efficiency up to the point where pre-TDC pressure work offsets the gains; in these conditions, IMEP peaked around 30° bTDC, with 35° improving efficiency but not work output. Beyond the specific numbers, the value lies in the reusable OpenFOAM infrastructure you built for OP engines and lean H_2 operation.

The DISI spray-model campaign, carried out in STAR-CCM+, complements the combustion chapters by pinning down practical choices that make Lagrangian spray simulations robust. Two findings stand out: (i) injector source points must be placed away from walls to avoid spurious jet collapse and unphysical wall interactions; and (ii) the chosen primary-breakup prescription (Blob or Rosin–Rammler) and the KHRT secondary-breakup tuning are coupled with the surrogate definition and chamber back-pressure. Under near-ambient back-pressure the Blob and Rosin–Rammler options were equivalent when matched on mean size; at elevated back-pressure, a larger initial size (with secondary-breakup doing the rest) produced realistic penetration and evaporation. Multi-component

surrogates required modest KHRT retuning due to parcel-averaged properties evolving as species evaporate at different rates. Importantly, across the tested operating conditions, the calibrated spray setup reproduced the experimental trends without modifying the previously defined calibration coefficients, underscoring its robustness and transferability.

Taken together, the chapters show that credible predictions come from aligning modeling assumptions with the operative regime, constraining models with targeted experiments or 1D references, and resisting unnecessary complexity. The premixed-charge approximation—validated by pre-cycle checks—enabled broad design sweeps at reasonable cost for ethanol; detailed chemistry clarified where simple flame models need help; open-source infrastructure opened the door to OP/H_2 concepts; and the spray calibration proved predictive across operating conditions without re-tuning calibration coefficients, turning disparate experimental inputs into a consistent nozzle-exit specification.

The work also makes transparent its limits and the most promising extensions. Premixed assumptions should be replaced by fully resolved mixture formation when wall films, stratification, or cycle-to-cycle variability become decisive. The OP/H_2 setup, with fresh/burnt tracers and custom volume/flux accounting, provides a basis for future multi-cycle simulations of scavenging and for exploring ignition-system sensitivity. The DISI setup invites validation against time-resolved spray imaging and extension to biofuels and efuels surrogates with temperature-dependent properties. Finally, several questions would benefit from higher-fidelity turbulence (e.g., LES) and from integrated system studies (hybridization, aftertreatment, WtW fuel impacts) to translate component-level improvements into real-world gains.

Overall**,** the thesis delivers not only case-specific results—E85 performance/emissions potential, E5-vs-E85 kinetics behavior, lean H_2 OP phasing, and a robust, predictive DISI spray setup—but also a disciplined CFD methodology that can be reused and extended across engines, fuels, and operating regimes.

References

[1] EIA projects global conventional vehicle fleet will peak in 2038 - U.S. Energy Information Administration (EIA) n.d. https://www.eia.gov/todayinenergy/detail.php?id=50096 (accessed September 10, 2025).

[2] Serrano JR, Novella R, Piqueras P. Why the Development of Internal Combustion Engines Is Still Necessary to Fight against Global Climate Change from the Perspective of Transportation. Appl Sci 2019;9:4597. https://doi.org/10.3390/app9214597.

[3] Yılbaşı Z. Biofuels, E-Fuels, and Waste-Derived Fuels: Advances, Challenges, and Future Directions. Sustainability 2025;17:6145. https://doi.org/10.3390/su17136145.

[4] Uchida N, Onorati A, Novella R, Agarwal AK, Abdul-Manan AF, Kulzer AC, et al. E-fuels in IC engines: A key solution for a future decarbonized transport. Int J Engine Res 2025:14680874251325296. https://doi.org/10.1177/14680874251325296.

[5] Recent Fuel-Based Advancements of Internal Combustion Engines: Status and Perspectives | Energy & Fuels n.d. https://pubs.acs.org/doi/10.1021/acs.energyfuels.5c00057?ref=pdf (accessed October 2, 2025).

[6] Dornoff J, Rodríguez F. Euro 7: The new emission standard for light- and heavy-duty vehicles in the European Union n.d.

[7] G. Anderson L. Ethanol fuel use in Brazil: air quality impacts. Energy Environ Sci 2009;2:1015–37. https://doi.org/10.1039/B906057J.

[8] Bae C, Kim J. Alternative fuels for internal combustion engines. Proc Combust Inst 2017;36:3389–413. https://doi.org/10.1016/j.proci.2016.09.009.

[9] Crowl DA, Jo Y-D. The hazards and risks of hydrogen. J Loss Prev Process Ind 2007;20:158–64. https://doi.org/10.1016/j.jlp.2007.02.002.

[10] Azeem N, Beatrice C, Vassallo A, Pesce F, Davide G, Guido C, et al. Comparative Analysis of Different Methodologies to Calculate Lambda (λ) Based on Extensive And systemic Experimentation on a Hydrogen Internal Combustion Engine. SAE Technical Paper; 2023. https://doi.org/10.4271/2023-01-0340.

[11] AVL. AVL FIRE M Useg Guide n.d.

[12] Wu Y, Pal P, Som S, Lu T. A SKELETAL CHEMICAL KINETIC MECHANISM FOR GASOLINE AND GASOLINE/ETHANOL BLEND SURROGATES FOR ENGINE CFD APPLICATIONS n.d.

[13] Sfriso S, Berni F, Fontanesi S, Paltrinieri S, Mortellaro FS, Rossi V, et al. A numerical framework for 3D Computational Fluid Dynamics (CFD) simulations of gasoline-fuelled burners to be adopted in catalyst preheating of vehicle exhaust systems. Appl Therm Eng 2025;263:125332. https://doi.org/10.1016/j.applthermaleng.2024.125332.

[14] Wu C-W, Chen R-H, Pu J-Y, Lin T-H. The influence of air–fuel ratio on engine performance and pollutant emission of an SI engine using ethanol–gasoline-blended

fuels. Atmos Environ 2004;38:7093–100. https://doi.org/10.1016/j.atmosenv.2004.01.058.

[15] Regner G, Johnson D, Koszewnik J, Dion E, Redon F, Fromm L. Modernizing the Opposed Piston, Two Stroke Engine for Clean, Efficient Transportation. SAE Technical Paper; 2013. https://doi.org/10.4271/2013-26-0114.

[16] Onorati A, Payri R, Vaglieco B, Agarwal A, Bae C, Bruneaux G, et al. The role of hydrogen for future internal combustion engines. Int J Engine Res 2022;23:529–40. https://doi.org/10.1177/14680874221081947.

[17] Weller HG, Tabor G, Gosman AD, Fureby C. Application of a flame-wrinkling les combustion model to a turbulent mixing layer. Symp Int Combust 1998;27:899–907. https://doi.org/10.1016/S0082-0784(98)80487-6.

[18] OpenFOAM. OpenFOAM User Guide n.d.

[19] Notes on Computational Fluid Dynamics: General Principles. CFD Direct 2022. https://doc.cfd.direct/notes/cfd-general-principles/index/ (accessed October 19, 2025).

[20] Stefano Fontanesi, Alessandro D'Adamo, Fabio Berni. Simulazione CFD dei motori a Combustione Interna n.d.

[21] Giuseppe Cantore. Motori Endotermici Alternativi con cenni di controllo. Prima. Società Editrice Esculapio; 2008.

[22] Impact of the Primary Break-Up Strategy on the Morphology of GDI Sprays in 3D-CFD Simulations of Multi-Hole Injectors n.d. https://www.mdpi.com/1996-1073/12/15/2890 (accessed October 20, 2025).

[23] SIEMENS. STAR-CCM+ User Guide n.d.

[24] Patterson MA, Reitz RD. Modeling the Effects of Fuel Spray Characteristics on Diesel Engine Combustion and Emission. SAE Trans 1998;107:27–43.