Politecnico di Torino

ScuDo
Scuola di Dottorato ~ Doctoral School
WHAT YOU ARE, TAKES YOU FAR

Doctoral Dissertation
Doctoral Program in Aerospace Engineering (38^{th} cycle)

Modelling of Direct Hydrogen Combustion in Jet Engines

System-Level Analysis, Computational Fluid Dynamics and Chemical Reactor Networks

By

Lorenzo Folcarelli

******

Supervisor(s):
Prof. Dario Giuseppe Pastrone, Supervisor
Prof. Andrea Ferrero, Co-Supervisor
Prof. Filippo Masseni, Co-Supervisor

Doctoral Examination Committee:
Prof. Manuela Battipede, Politecnico di Torino
Prof. Pietro Paolo Ciottoli, Referee, Sapienza Università di Roma
Prof. Maria Grazia De Giorgi, Referee, Università del Salento

Politecnico di Torino
2026

Declaration

I hereby declare that, the contents and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

Lorenzo Folcarelli
2026

* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Politecnico di Torino (ScuDo).

Alla mia famiglia.
Ad Alessia.

Acknowledgements

I would like to express my deepest gratitude to my supervisors for their invaluable guidance, each of whom has left a distinct and lasting mark on my journey. I am profoundly grateful to Prof. Dario Pastrone, who has been an unwavering cornerstone of support, teaching me essential values that extend far beyond mere professional development. I sincerely thank Prof. Andrea Ferrero for opening the doors to industrial and international opportunities through his network, and for providing precise, insightful feedback that consistently elevated the quality of my research. My heartfelt thanks also go to Prof. Filippo Masseni, who accompanied me to my first conferences and, through his trust, gave me the confidence to truly believe in my research project.

I am grateful to Dr. Hysam Telib for hosting me at Optimad, and to Anna Spagnolo and Dr. Fausto Dicech for their active support, outstanding availability, and for everything they taught me during the study of the AHEAD combustor. In particular, I am deeply indebted to them for granting me access to their multi-fidelity model, which I effectively employed in this thesis, and I gratefully acknowledge this valuable contribution to my work.

I would like to thank Prof. Fabio Cozzi for his availability and kindness in providing detailed data on his combustor configuration, and for taking the time to answer my questions and offer insights regarding his experimental setup.

My time abroad was deeply transformative, and I owe a great debt of gratitude to Prof. Markus Klein. Through his immense dedication and commitment, he acted as a true supervisor during my six months at the Universität der Bundeswehr München, not only providing access to his computational resources but, above all, sharing his extensive experience and passion for combustion modelling. Alongside him, I am deeply grateful to Dr. Josef Hasslberger for his passionate guidance; his constant readiness to delve deeper, clarify, and explain highlighted an incredible

work ethic and an extraordinary human quality. At UniBw, I was also fortunate to meet Dr. Riccardo Concetti, whom I thank for becoming a true friend and for collaboratively developing the chemical reactor network methodology.

I wish to extend my deepest appreciation to all the colleagues who shared this journey with me and made these years so much lighter. A special thanks goes to my office mates — Giacomo Gedda, Giovanni Polizzi, Leonardo Stumpo, Alessandra Zumbo, Vittorio Delnegro, Vincenzo Madonia, Sebastiano Marcì, Daniele Tozzi, Isabelle Veith and Wassim Cherkaoui — for the daily support, laughs, and camaraderie we shared. Among all of them, I am uniquely indebted to Giacomo; we began our university path together on our very first days in Torino, and walking every single step of this long academic journey side by side, all the way through the PhD, has been a rare and constant blessing. I am deeply grateful to Vincenzo for his relentless dedication and hard work. Sharing the same research focus made the challenges much more rewarding, and his constant support was key to expanding and refining this thesis, helping us both grow along the way. Finally, I would like to thank Sebastiano for his valuable help during the final stages of this work.

On a personal note, I want to thank my family, whose unconditional support allowed me to face this challenging path with peace of mind and full focus on my studies. A special, silent thought goes to my grandparents; though they are no longer here to share this milestone, their guidance shaped the person I am today. My deepest, most affectionate, and immense gratitude goes to my girlfriend, Alessia. We met just before I began this journey, and she has been my constant pillar and my safe harbour ever since. These years have been a beautiful journey of growing together, and her unwavering trust in me has been an incredible source of strength. Even when the heavy demands of this thesis left us with so little time, she brought joy into my darkest moments, helping me not only to work better, but above all, to live better. I am profoundly lucky to have her by my side, and to share my life and this achievement with her.

Finally, I wish to honour the memory of Dr. Filipe S. Pereira. Although our path crossed only for a short time, he willingly shared his knowledge with a young researcher. His early advice on uncertainty quantification significantly influenced my approach to this work. I am deeply grateful for his kindness, his encouragement, and the guidance he so freely offered.

This work was supported by two institutions whose contributions I gratefully acknowledge: the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time on the GCS Supercomputer, and HPC@PoliTO (www.hpc.polito.it) for providing computational resources at Politecnico di Torino. This publication is part of the project PNRR-NGEU, which has received funding from the Italian Ministry of University and Research (MUR) (Ministerial Decree 351/2022).

Abstract

The decarbonization of aviation requires a fundamental shift in propulsion technologies, with hydrogen emerging as a promising zero-carbon fuel for gas turbine engines. Although hydrogen possesses unique combustion characteristics, such as high flame speeds and wide flammability limits, it also presents susceptibility to flashback and the potential for NO_x production due to high combustion temperatures, posing significant design challenges that demand efficient computational tools for practical engineering applications. This thesis develops and validates computational methodologies that enable accurate and efficient analysis of hydrogen combustion in jet engines, addressing the critical challenges of flashback prevention and nitrogen oxide emission control.

Research progresses along three complementary fronts. First, system-level thermodynamic analysis using the Numerical Propulsion System Simulation (NPSS) software demonstrates the feasibility and performance benefits of hydrogen-fuelled turbofan cycles. By modelling the complete cryogenic fuel system from tank to combustor, this work evaluates alternative thermodynamic configurations exploiting hydrogen's heat sink capacity, revealing potential improvements in cycle efficiency.

Second, Reynolds-Averaged Navier-Stokes (RANS) simulations are employed to characterize hydrogen flame behaviour in two configurations: a micromix combustor with converging nozzle geometries and the AHEAD lean premixed swirl-stabilized burner featuring radial swirl generation and axial air injection for flashback mitigation, investigated with both two-dimensional axisymmetric and three-dimensional RANS. A multi-fidelity modelling framework is used to integrate computationally inexpensive two-dimensional axisymmetric RANS simulations with limited experimental data for flame position prediction. This approach achieves accurate flame position assessment with a number of calibration points comparable to the size of the

input parameter space, outperforming more complex surrogate modelling techniques at small training set sizes.

Third, representing the core contribution of this work, a Chemical Reactor Network (CRN) methodology is developed for rapid prediction of NO_x emissions without requiring the full resolution of the combustor flow field. The methodology employs an optimization algorithm for network calibration against high-fidelity data and incorporates physically motivated reactor topologies derived from fundamental swirl combustion structures. Validation on a non-premixed hydrogen combustor with water injection demonstrates the model's capability to capture the exponential NO_x reduction achieved through steam dilution. Application to the AHEAD premixed combustor shows that the calibrated network successfully predicts emission trends across wide operating ranges. A comparison against two-dimensional and three-dimensional RANS demonstrates that the CRN achieves accuracy competitive with or superior to RANS simulations, with a cost separation exceeding five orders of magnitude against 3D RANS. The robustness of these predictions is further assessed through a Polynomial Chaos Expansion (PCE) uncertainty analysis, which identifies inlet flow-rate measurement accuracy as the dominant contributor to NO_x prediction variance, rather than the calibrated model parameters. In both test cases, the CRN exhibits the ability to generalize results far from the conditions used for its calibration, enabling design space exploration that would otherwise be computationally prohibitive.

The computational tools and methodologies developed in this thesis provide a foundation for emissions-conscious preliminary design of hydrogen aviation propulsion systems, balancing accuracy, physical insight, and computational tractability, essential for industrial design processes.

Contents

List of Figures

List of Tables

Nomenclature

Roman Symbols

A	Area [m^2]
$a_{i,j}$	NASA7 polynomial coefficient for species j [-]
C	Heat capacity rate [W/K]
C_0	Root chord length [m]
$CI_{95}^{lo}, CI_{95}^{hi}$	lower and upper bounds of the 95% uncertainty interval [ppm]
c_p	Specific heat at constant pressure [J/(kg·K)]
$\hat{c}_{p,j}$	Molar heat capacity at constant pressure of species j [J/(mol·K)]
C_τ	EDC time scale constant [-]
C_ξ	EDC model constant [-]
D	Diameter of the mixing tube of the AHEAD combustor [m]
d	PCE input dimensionality [-]
D_m	Mixture diffusivity [m^2/s]
Da	Damköhler number [-]
D_i	Inner diameter [m]
$D_{j,k}$	Generalized Fick's law diffusion coefficients [m^2/s]
d_min	Minimum distance to training points [-]

D_o Outer diameter [m]

ΔT Temperature difference [K]

$D_{\tilde{T},j}$ Soret diffusion coefficient of species j [kg/(m·s)]

ΔT_t Total temperature drop [K]

d_w Wall distance [m]

E Specific total energy (internal + kinetic) [J/kg]

E_{LD} Lift-to-drag ratio [-]

e Specific internal Energy [J/kg]

F Thrust [N]

f Objective function [-]

f_{hy} WMLES blending function [-]

$f_{\mathrm{st,H_2/O_2}}$ Stoichiometric mass ratio of hydrogen to oxygen

g Gravitational acceleration [m/s^2]

h Specific sensible enthalpy [J/kg]

H Specific total enthalpy [J/kg]

H_{ref} Reference fuselage height [m]

H_{stage} Head per stage [ft]

$\vec{J}_j$ Diffusion flux of species j [kg/(m^2·s)]

k Laminar thermal conductivity [W/(m·K)]

k_{eff} Effective thermal conductivity [W/(m·K)]

k_i Thermal conductivity of layer i [W/(m·K)]

L_{tank} Tank length [m]

M Total number of species in mixture [-]

m	Mass [kg]
$\dot{m}$	Mass flow rate [kg/s]
$\mathscr{M}_j$	Molecular weight of species j [kg/mol]
N	Rotational speed [rpm]
n	Number of high-fidelity training points [-]
N_{ext}	Size of the independent external validation set [-]
N_{reg}	PCE regression sample size per equivalence-ratio set point [-]
N_s	Pump specific speed [$\text{rpm}\,\text{gpm}^{1/2}\,\text{ft}^{-3/4}$]
p	Pressure [Pa]
p_{t}	Total pressure [kPa]
P_{basis}	Number of PCE basis functions [-]
p_o	PCE polynomial order [-]
$\mathscr{P}$	Optimization parameter in the non-premixed test case [-]
Pr_t	Turbulent Prandtl number [-]
$\dot{Q}$	Heat transfer rate [W]
Q^2_{LOO}	Stone–Geisser leave-one-out predictive coefficient, [–]
Q_{p}	Pump volumetric flow rate [gpm]
R	Universal gas constant [J/(mol·K)]
r	Radial coordinate [m]
R_b	Annular air passage [m]
R_{cc}	Combustion chamber radius
$\mathscr{R}_j$	Volumetric rate of creation of species j [kg/(m^3·s)]
R_{mix}	Mixing tube radius

R^2_{test} Coefficient of determination on the external test set, [–]

R_{th} Specific thermal resistance [$m^2 \cdot$ K/W]

Sc_t Turbulent Schmidt number [-]

S_{ij} Strain rate tensor [1/s]

S_L Laminar flame speed [m/s]

S_{NO} NO chemical source term [kg/($m^3 \cdot$s)]

S_{NO_2} NO_2 chemical source term [kg/($m^3 \cdot$s)]

S Swirl number [-]

S_i^T Total-effect Sobol index of uncertain input i [–]

$\langle S_T \rangle_\phi$ ϕ-averaged total-effect Sobol index [–]

T Temperature [K]

T_{t} Total temperature [K]

t_{Al} Aluminium layer thickness [m]

t_{CFRP} CFRP layer thickness [m]

t_{foam} Foam insulation thickness [m]

t_i Layer thickness [m]

t Time [s]

T_u Unburned mixture temperature [K]

$\vec{v}$ Velocity vector [m/s]

V Volume [m^3]

V_{LH_2} Liquid hydrogen volume [m^3]

v_r Radial velocity component [m/s]

V_{tank} Tank volume [m^3]

v_θ Tangential (swirl) velocity component [m/s]

v_x Axial velocity component [m/s]

w Weight function [-]

w_t Turbine specific work [J/kg]

X Mole Fraction [-]

x Axial coordinate [m]

Y Mass fraction [-]

y^+ Dimensionless wall distance [-]

Z Elemental mass fraction

Greek Symbols

β Generic scalar

γ Deflection angle of swirled flow [rad]

δ_a, δ_f relative deviations of air and fuel mass-flow rates [–]

δ_{ij} Kronecker delta [-]

δ Correction field [-]

ε Rate of dissipation of turbulent kinetic energy [m^2/s^3]

ε_{HX} Heat exchanger effectiveness [-]

η_{grav} Gravimetric efficiency [-]

θ_k k-th optimization parameter in the premixed test case [-]

$\vec{\Theta}$ Optimization vector in the premixed test case [-]

$\vec{\Theta}^*$ reference calibrated parameter vector in the premixed test case [–]

θ Azimuthal coordinate [rad]

κ Turbulent kinetic energy [m^2/s^2]

κ_{PaSR} PaSR reacting volume fraction [-]

κ_{fs} EDC fine-scale reaction-rate pre-factor [-]

μ Dynamic viscosity [Pa·s]

μ_{NO_x} PCE-predicted mean NO_x, ppm

μ_{sgs} Subgrid-scale dynamic viscosity [Pa·s]

μ_t Turbulent dynamic viscosity [Pa·s]

ν Kinematic viscosity [m^2/s]

ξ Fine-structure region mass fraction [-]

ϕ Generic filtered scalar

ρ Density [kg/m^3]

σ_{NO_x} PCE-predicted standard deviation of NO_x [ppm]

$\hat{\tau}$ Total stress tensor [Pa]

τ^* EDC time scale [s]

τ_c Chemistry time scale [s]

τ_{mix} Mixing time scale [s]

τ_{ij}^{sgs} Subgrid-scale stress tensor [Pa]

ϕ Equivalence ratio [-]

χ Axial injection ratio [-]

Ω Volumetric heat of reaction [W/m^3]

ω Specific dissipation rate [1/s]

Superscripts

$(\cdot)^*$ Fine-scale value

$(\cdot)^0$ Formation property

$(\cdot)^{opt}$ Optimization

Subscripts

$(\cdot)_{axial}$ Axial component

$(\cdot)_{dil}$ Dilution air

$(\cdot)_{HF}$ High-fidelity value

$(\cdot)_{in}$ Inlet

$(\cdot)_j$ Property of species j

$(\cdot)_{LF}$ Low-fidelity value

$(\cdot)_{out}$ Outlet

$(\cdot)_{ref}$ Reference

$(\cdot)_{swirl}$ Swirl component

$(\cdot)_{tot}$ Total

Other Symbols

$\overline{(\cdot)}$ Reynolds-averaged value

$\nabla\cdot$ Divergence operator [1/m]

∇ Gradient operator [1/m]

$\widetilde{(\cdot)}$ Favre-averaged value

Acronyms / Abbreviations

AHEAD Advanced Hybrid Engines for Aircraft Development

AIT Autoignition temperature

BPR Bypass ratio

BWB Blended Wing Body

Cantera Open-source chemical kinetics software

CCS Carbon capture and storage

CEA Chemical Equilibrium with Applications

CFD Computational Fluid Dynamics

CFL Courant–Friedrichs–Lewy condition

CFRP Carbon fibre reinforced plastic

CIVB Combustion-induced vortex breakdown

CMA-ES Covariance Matrix Adaptation Evolution Strategy

CO_2 Carbon dioxide

CoolProp Open-source thermodynamic properties library

CRECK Chemical REaction Engineering and Chemical Kinetics

CRN Chemical reactor network

EARSM Explicit Algebraic Reynolds Stress Model

EASA European Union Aviation Safety Agency

EDC Eddy Dissipation Concept

ENOVAL Environmentally Optimized Turbofan

FAA Federal Aviation Administration

FPR Fan pressure ratio

GH_2 Gaseous hydrogen

HySITE Hydrogen Steam and Inter-Cooled Turbine Engine

ICAO International Civil Aviation Organization

LES Large Eddy Simulation

LH_2 Liquid hydrogen

LHV Lower heating value

MAE Mean Absolute Error

MAPE Mean Absolute Percentage Error

MF-BWB Multi-Fuel Blended Wing Body

MIE Minimum ignition energy

MILD Moderate or Intense Low-oxygen Dilution

NACA National Advisory Committee for Aeronautics

NARGP Nonlinear AutoRegressive Gaussian Process

NO_x Nitrogen oxides

NPSS Numerical Propulsion System Simulation

NTC Negative temperature coefficient

NTU Number of transfer units

OPR Overall pressure ratio

PaSR Partially Stirred Reactor

PCE Polynomial Chaos Expansion

PDF Probability Density Function

PIV Particle Image Velocimetry

PLIF Planar Laser-Induced Fluorescence

PSO Particle Swarm Optimization

RANS Reynolds-Averaged Navier-Stokes

RBF Radial Basis Function

RMSE Root Mean Square Error

SAFs Sustainable aviation fuels

SGS Subgrid-Scale

SimpleMF Simplified Multi-Fidelity

SST Shear Stress Transport

TIT Turbine inlet temperature

TSEC Thrust-specific energy consumption

TSFC Thrust-specific fuel consumption

UHBPR Ultra-high bypass ratio

WALE Wall-Adapting Local Eddy-viscosity

WMLES Wall-Modeled Large Eddy Simulation

Chapter 1

Introduction

The global imperative to mitigate climate change has positioned the aviation sector at a critical juncture. As one of the fastest-growing contributors to greenhouse gas emissions, aviation faces unprecedented pressure to decarbonize while meeting the increasing demand for air transportation. This thesis investigates direct hydrogen combustion as a pathway towards sustainable jet propulsion, assessing low-cost CFD methods for the analysis of hydrogen flames and proposing a new methodology for NO_x prediction through chemical reactor networks.

1.1 The Need for Sustainable Aviation

The current aviation industry contributes approximately 3.5% of global greenhouse gas emissions produced by human activities, and this percentage is expected to increase substantially in the coming decades [83]. Despite ongoing efforts by aircraft manufacturers to reduce fuel consumption per seat mile through improved aerodynamics, lightweight materials, and advanced engine technologies, the growing demand for air transportation is predicted to outpace these technological advancements in fuel efficiency [2]. Projections indicate that without transformative changes to propulsion systems, aviation's contribution to climate change will become increasingly heavier.

In response to this challenge, the European Commission has established ambitious targets to mitigate the climate impact of aviation. The Flightpath 2050 initiative aims for a 75% reduction of carbon dioxide (CO_2) emissions and a 90% reduction of

nitrogen oxide (NO_x) emissions by 2050, compared to the capabilities of a typical new aircraft in 2000 [27]. Similarly, the International Civil Aviation Organization (ICAO) has set goals for carbon-neutral growth beyond 2020 and a 50% reduction in net aviation CO_2 emissions by 2050 relative to 2005 levels. These targets represent a fundamental departure from incremental improvements and necessitate the exploration of alternative propulsion concepts.

Among the various pathways towards sustainable aviation, including sustainable aviation fuels (SAFs), electric propulsion, and hybrid-electric systems, hydrogen-fuelled gas turbine engines emerge as one of the most promising solutions for medium- to long-range aircraft [2]. Unlike drop-in SAFs, which can reduce lifecycle emissions but still produce CO_2 during combustion, hydrogen offers the potential for zero direct carbon emissions while leveraging the well-established gas turbine architecture that has proven reliable over decades of commercial aviation. This compatibility with existing propulsion technology, combined with hydrogen's properties, positions direct hydrogen combustion as a compelling option for decarbonizing aviation.

1.2 Hydrogen as an Aviation Fuel

Hydrogen possesses a unique combination of properties that make it both highly attractive and technically challenging as an aviation fuel. Understanding these fundamental characteristics is essential for appreciating the design considerations that permeate this research.

1.2.1 Fundamental Properties of Hydrogen

Hydrogen exhibits several distinctive physical and chemical properties that differentiate it from conventional jet fuels such as Jet-A or kerosene. The most significant advantage of hydrogen lies in its exceptionally high gravimetric energy density. The lower heating value (LHV) of hydrogen is approximately 120 MJ/kg, compared to 43 MJ/kg for kerosene. This 2.8-times higher energy content per unit mass means that, in principle, an aircraft would require significantly less fuel mass to complete a given mission, potentially reducing the aircraft's structural weight and improving overall efficiency [56].

However, this advantage is substantially offset by hydrogen's extremely low volumetric energy density. At atmospheric pressure and ambient temperature, hydrogen is a gas with negligible energy density per unit volume. Even when liquefied at cryogenic temperatures (approximately 20 K at atmospheric pressure), liquid hydrogen (LH_2) has an energy density per unit volume about four times lower than that of kerosene. This fundamental mismatch between gravimetric and volumetric energy density creates a central design challenge for hydrogen aircraft: the fuel tanks must be significantly larger than those of conventional aircraft, potentially introducing weight penalties from the tank structure itself and increased aerodynamic drag from accommodating the bulky fuel system [2].

The combustion characteristics of hydrogen also differ markedly from those of kerosene. Hydrogen exhibits a much wider flammability range [21], e.g. it has flammability limits in air ranging from approximately 4% to 75% by volume [127], compared to 1% to 5% for kerosene vapor [31], at standard conditions. This broad flammability range offers operational advantages, including enhanced ignition reliability and the ability to operate at very lean equivalence ratios, which is beneficial for reducing pollutant emissions. Furthermore, hydrogen's wide flammability range enhances the engine's capability to reignite in case of an emergency, providing an important safety advantage in critical flight situations.

Another critical property is hydrogen's high laminar flame speed, which at stoichiometric conditions and standard temperature and pressure is approximately 2.5 m/s, compared to 0.4 m/s for typical hydrocarbon fuels [79]. This high flame speed, combined with hydrogen's elevated mass diffusivity (approximately 0.61 cm^2/s at 298 K, compared to 0.05 cm^2/s for typical hydrocarbons), significantly increases the propensity for flame flashback in premixed combustion systems [17]. Flashback, wherein the flame propagates upstream into the fuel-air mixing region, poses severe safety risks including hardware damage and potential catastrophic failure, making its prevention a paramount concern in hydrogen combustor design.

The adiabatic flame temperature of hydrogen-air mixtures provides further insight into combustion behaviour. Figure 1.1 presents a comparison between hydrogen-air and Jet-A-air adiabatic flame temperatures as functions of equivalence ratio, calculated using NASA Chemical Equilibrium with Applications (CEA) for representative combustor inlet conditions (fuel at 300 K, air at 600 K, pressure of 6.9 bar) [51]. At stoichiometric conditions ($\phi = 1$), the adiabatic flame temperature of the H_2-air mix-

ture is approximately 4% higher than that of the Jet-A-air mixture, reaching values around 2400 K. The peak flame temperature occurs slightly rich of stoichiometric due to dissociation effects [82]. These elevated temperatures directly influence the formation of nitrogen oxides, which constitute the primary pollutant in hydrogen combustion with air, as discussed further in Section 1.5.2.

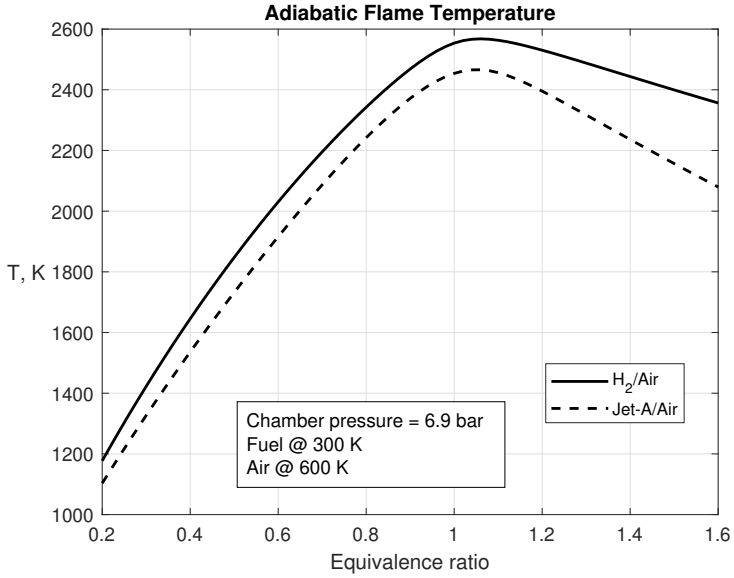

Fig. 1.1 Comparison between hydrogen-air and Jet-A-air adiabatic flame temperatures as functions of equivalence ratio. Conditions: fuel at 300 K, air at 600 K, pressure 6.9 bar. Data calculated using NASA CEA.

1.2.2 Production Methods

The environmental credentials of hydrogen as an aviation fuel are intrinsically linked to its production pathway. Hydrogen production is commonly classified by a colour-coding system that reflects the carbon intensity of the production process.

Green hydrogen is produced through the electrolysis of water using renewable electricity sources such as wind, solar, or hydroelectric power. This pathway offers the potential for truly zero-carbon hydrogen, making it the most environmentally desirable option for sustainable aviation. However, green hydrogen currently faces challenges related to cost, energy efficiency, and the availability of sufficient renewable electricity capacity. The round-trip energy efficiency of producing hydrogen via electrolysis, storing it, and then converting it back to useful work in a fuel cell or

combustion engine is typically lower than direct use of electricity, which limits the overall system efficiency.

Blue hydrogen is produced through steam methane reforming of natural gas, with the resulting CO_2 captured and sequestered through carbon capture and storage (CCS) technologies. While not entirely carbon-free due to fugitive emissions and the energy required for the capture process, blue hydrogen represents a potentially lower-cost near-term option with substantially reduced emissions compared to conventional fossil fuel combustion. The maturity of steam methane reforming technology and existing natural gas infrastructure provide a pathway for scaling up hydrogen production more rapidly than green hydrogen alone.

Other hydrogen production pathways include gray hydrogen (steam methane reforming without CCS), brown hydrogen (from coal gasification), and emerging technologies such as thermochemical water splitting and biological production. For the purposes of this thesis, the focus is on the combustion behaviour and system integration of hydrogen regardless of its production pathway, with the understanding that achieving net-zero aviation emissions ultimately requires a transition to green hydrogen production at scale.

1.2.3 Storage Technologies

The storage of hydrogen aboard an aircraft represents one of the most significant technical challenges in hydrogen aviation. Two primary storage approaches have been considered: compressed gaseous hydrogen and liquid hydrogen.

Compressed gaseous hydrogen storage involves maintaining hydrogen at high pressures, typically 350 to 700 bar, in thick-walled pressure vessels. While this approach simplifies the thermal management system by eliminating the need for cryogenic cooling, it suffers from extremely poor volumetric energy density even at high pressures. The resulting tank weight and volume penalties make compressed hydrogen storage impractical for all but the smallest aircraft, such as regional commuter planes or urban air mobility vehicles. For medium- to long-range transport aircraft, the size and weight of high-pressure tanks would be prohibitive [2].

Liquid hydrogen storage, despite its complexity, is generally considered the only viable option for larger commercial aircraft. By cooling hydrogen to approximately 20 K at atmospheric pressure, the density increases to about 71 kg/m^3, significantly

improving the volumetric energy density compared to gaseous storage. However, even in liquid form, LH_2 requires approximately four times the volume of kerosene for equivalent energy content [56].

The design of LH_2 tanks for aircraft involves several critical considerations. First, the tanks must be heavily insulated to minimize heat ingress from the ambient environment, which would cause boil-off (the evaporation of liquid hydrogen due to heat absorption). Multi-layer insulation systems, typically consisting of vacuum-insulated walls with layers of aluminium foil and foam, are employed to achieve thermal resistance sufficient to limit boil-off rates to acceptable levels, typically targeting less than 2% of the stored fuel mass over a mission duration [158]. Second, the tanks must be structurally sound to withstand the combination of internal pressure (which increases as boil-off occurs if the tank is not vented), external aerodynamic loads, and mechanical loads during flight maneuvers. Third, safety considerations require robust containment systems to prevent hydrogen leakage, which poses both explosion hazards and environmental concerns due to hydrogen's tendency to embrittle certain materials.

Tank configuration is closely coupled with aircraft design. For conventional tube-and-wing aircraft, integrating large cylindrical LH_2 tanks into the fuselage presents significant challenges, often requiring relocation of passengers or cargo. This has motivated interest in alternative aircraft configurations, such as the Blended Wing Body (BWB) architecture, which offers larger internal volumes that can more readily accommodate hydrogen tanks while minimizing the aerodynamic penalties associated with external storage. The CRYOPLANE project and subsequent studies have explored various tank placement strategies and their implications for aircraft centre of gravity, structural loads, and overall system mass [157].

Boil-off management is another critical aspect of LH_2 storage. The vaporized hydrogen from boil-off can potentially be used beneficially, for example by routing it through heat exchangers to precool or intercool compressor air, thereby improving thermodynamic cycle efficiency. However, careful system design is required to ensure that boil-off rates remain predictable and manageable across all phases of flight and ground operations. Extended ground hold times before takeoff, in particular, can lead to significant boil-off, necessitating venting or reliquefaction systems.

The gravimetric efficiency of the fuel storage system, defined as the ratio of stored hydrogen mass to the total mass of the fuel system (including tanks, insulation, and associated hardware), is a key performance metric. Achieving high gravimetric efficiencies is generally considered necessary for practical hydrogen aircraft, requiring advanced materials and optimized structural designs [2].

1.3 Historical Use of Hydrogen in Aviation

The concept of hydrogen-fuelled aircraft is not new; it has been explored intermittently over the past several decades, with varying degrees of technical success and programmatic momentum. This historical context provides important lessons for contemporary efforts to develop hydrogen aircraft technologies.

Hydrogen had long been used in lighter-than-air flight in balloons and airships. Remarkably, its first application in jet propulsion coincided with the birth of the jet engine itself: von Ohain's HeS 1 demonstrator tested in March 1937 at Heinkel's Marienehe works [101]. Von Ohain chose gaseous hydrogen for this first demonstrator to facilitate rapid testing and meet Ernst Heinkel's demanding one-year development timeline. The decision proved sound: the hydrogen combustor performed reliably due to hydrogen's high flame velocity and wide combustion range, operating without significant problems even under off-design conditions and during transient acceleration and deceleration [101]. This successful demonstration in spring 1937 validated the fundamental turbojet concept. In parallel, von Ohain's team initiated investigations into gasoline combustion systems, recognizing that operational engines would require hydrocarbon fuels. This development pathway led to HeS 3 engines, which utilized gasoline and culminated in the historic first jet-powered flight of the He 178 on August 27, 1939 [101].

After that initial attempt, interest in liquid hydrogen for aircraft propulsion emerged in the United States in 1955, when Silverstein and Hall of the then NACA-Lewis Flight Propulsion Laboratory published a report on the potential of liquid hydrogen as a fuel in terms of range and flight altitude [131]. Based on this report, an experimental campaign was launched on the use of liquid hydrogen at high altitude by modifying a U.S. Air Force B-57 twin engine. The aircraft took off and reached cruise altitude using kerosene, and then used hydrogen, gasified in a heat exchanger, in one of the two engines for about 11 minutes, before reverting back to kerosene

once hydrogen was depleted. Three successful flights were made, and on each occasion the engine operated without incidents [36].

In 1956, the U.S. Air Force contracted Lockheed's Advanced Development Projects to build two CL-400 prototype aircraft for high-altitude reconnaissance, designed to cruise at Mach 2.5 at 100,000 feet with a range of 2200 nmi. Concurrently, Pratt & Whitney was awarded a contract to investigate the feasibility of using liquid hydrogen as aircraft fuel. A J57 turbojet was converted to operate on LH_2 in just 5 months, showing that the conversion of the engine to liquid hydrogen was "far easier than one might imagine", and tested in fall 1956, demonstrating that acceptable combustion efficiency could be achieved with combustor cans as short as one fourth the length required for conventional hydrocarbon fuels [105]. A new engine design, the Model 304, was subsequently developed using the expander cycle with heat exchanger, accumulating 25 hours of testing under sea level conditions. Despite successfully demonstrating the technical feasibility of hydrogen-fuelled flight, the CL-400 program was terminated in 1957. The design suffered from the combination of a relatively small wing and the large fuselage volume required to contain the low-density LH_2, while the Model 304 engine, though a technical achievement, proved to be heavy and relatively inefficient. Additionally, the logistics problem of requiring worldwide LH_2 availability for reconnaissance missions made the aircraft impractical for its intended role [23]. Nevertheless, the expertise gained from these pioneering projects proved invaluable for the United States space program, where hydrogen propulsion found its most successful application in rocket engines.

Following the hiatus in aircraft-focused hydrogen research during the 1960s, interest was renewed by the 1973 oil crisis. The most comprehensive early study of hydrogen aviation was conducted by Lockheed under NASA sponsorship from 1973 to 1978, with emphasis shifting from high-altitude reconnaissance to commercial transport applications. This work developed detailed fuel system designs for a 400-passenger subsonic transport with a 10,190 km range at Mach 0.85 cruise speed [22]. The study evaluated numerous configurations including fuel stored in fuselage tanks, wing-mounted pods, and external tanks on pylons above the wing. The study systematically evaluated fuel system architectures including boost pump designs, feed line insulation concepts, high-pressure fuel delivery systems and five thermodynamic cycles exploiting hydrogen's cryogenic properties: compressor air precooling, compressor intercooling, hydrogen cooling of turbine cooling air, regenerative exhaust gas heating of the fuel, and an expander cycle using hydrogen

expansion turbines for accessory power. The regenerative exhaust heating cycle, employing counterflow heat exchangers with the core exhaust stream, yielded the largest performance benefit, a 4.3% reduction in specific fuel consumption, and was selected as the preferred concept. Additionally, the fuel's heat sink capacity was exploited through multiple heat exchangers to cool turbine cooling air, engine oil, and environmental control system air, eliminating the need for conventional mechanical refrigeration. This work was later expanded by Brewer in the text "Hydrogen Aircraft Technology" [23].

Between 1988 and 1989, the Soviet Union conducted flight tests of the Tupolev Tu-155, an experimental aircraft based on the Tu-154 airliner with one of its three NK-8-2U turbofan engines converted to the experimental NK-88 configuration capable of operating on either liquid hydrogen or liquefied natural gas [137]. The first flight on liquid hydrogen occurred on April 15, 1988, followed by the first liquefied natural gas flight on January 18, 1989. Throughout the test campaign, the modified engine operated continuously on cryogenic fuel during all phases of flight, including takeoff, cruise, and landing. No incidents or hydrogen-system failures were reported [137]. These successful demonstrations corroborated the technical feasibility of hydrogen-fuelled aircraft, reinforcing the conclusion previously established by Lockheed's research program in the 1970s [22].

Building upon these early experiences, the European Commission has funded a series of comprehensive projects over the past two decades aimed at systematically addressing the multifaceted challenges of hydrogen aviation as a pathway towards decarbonization. The first was the CRYOPLANE project (2000-2003), which involved 36 partner organizations in a systematic examination of the feasibility of liquid hydrogen in all aspects of commercial aircraft operations [157]. The project developed detailed conceptual designs for seven aircraft categories spanning business jets to very large long-range aircraft (A380-class), establishing three fundamental tank layout strategies: aft fuselage tanks for smaller aircraft, combined aft and dorsal fuselage tanks for regional and short/medium-range aircraft, and fore and aft fuselage tanks for long-range aircraft. Propulsion system studies confirmed that conventional turbofan engines could be converted to hydrogen operation through combustor redesign and addition of fuel vaporization heat exchangers, with the engines operating 30-50 K cooler than kerosene equivalents at the same thrust level; ground testing of a "micromix" hydrogen combustor in an auxiliary power unit demonstrated NO_x emission reductions of up to 80% compared to kerosene com-

bustion [157]. Comprehensive safety assessments concluded that hydrogen-fuelled aircraft could achieve equivalent safety levels to kerosene aircraft through appropriate design measures and operational procedures, though existing airworthiness regulations would require adaptation. The project identified no fundamental technical barriers to implementation but emphasized the need for extensive component development and infrastructure investment.

The Advanced Hybrid Engines for Aircraft Development (AHEAD) project (2011-2015) was also funded by the European Union. The project developed an integrated approach combining aircraft configuration, propulsion system design, and climate impact assessment for a Multi-Fuel Blended Wing Body (MF-BWB) aircraft powered by novel hybrid engines [4]. The hybrid engine architecture featured dual combustion chambers: a lean premixed swirl-stabilized primary combustor for liquid hydrogen or liquefied natural gas, and a secondary flameless combustor operating on biofuel or kerosene between turbine stages. This configuration exploited hydrogen's cryogenic properties for turbine cooling air conditioning while addressing the challenge of flashback through axial air injection, demonstrating stable operation without flashback incidents across the operational envelope [4]. The hydrogen combustor achieved NO_x emissions in the single-digit parts-per-million range at cruise conditions, while the flameless combustor demonstrated similarly low emissions under vitiated air conditions [4]. Performance analyses indicated that the MF-BWB aircraft concept could achieve reductions exceeding 50% in CO_2 emissions and 80% in NO_x, soot, and CO emissions compared to conventional aircraft of similar capacity. Climate modelling revealed that contrail formation characteristics differed significantly from kerosene-fuelled aircraft due to the altered water vapour emission indices, though reduced soot emissions were found to decrease contrail optical depth and radiative forcing [4]. The project's systematic approach, encompassing aircraft design, engine development, combustion technology demonstration, and climate assessment, provides important insights for contemporary efforts to develop sustainable hydrogen aircraft systems.

Building upon these earlier research efforts, the European Commission-funded ENABLEH2 project (2018-2022) advanced hydrogen aviation technology with particular emphasis on ultra-low NO_x combustion systems and fuel system thermal management [34]. The project matured micromix diffusion combustion technology, which exploits hydrogen's wide flammability limits and high molecular diffusivity to achieve lean combustion without the auto-ignition and flashback risks associated

with premixed systems. Through combined numerical and experimental investigations on injector arrays, full annular combustor segments, and altitude-relight studies, the project demonstrated that low momentum flux ratio injector designs delivered the lowest NO_x emissions while exhibiting reduced susceptibility to thermoacoustic instabilities compared to conventional kerosene combustion systems [34]. Performance assessments indicated that liquid hydrogen-fuelled aircraft could achieve 40-60% reductions in mission NO_x emissions relative to kerosene or sustainable aviation fuel counterparts. The project investigated hydrogen's heat sink potential for advanced propulsion technologies, concluding that while existing engine surfaces could preheat fuel with high heat transfer rates, their limited area restricted performance benefits [34]. ENABLEH2 developed comprehensive safety guidelines for liquid hydrogen implementation at aircraft, airport, and operational levels, supported by experimental determination of flammability limits and burning velocities across relevant temperature and pressure ranges. Technology evaluation studies established that multiple innovation waves would be necessary for hydrogen aircraft entry into service, with initial focus on maturing key technologies and establishing certification rules, followed by subsequent waves exploiting advanced concepts such as turboelectric distributed propulsion [34]. The project concluded that hydrogen-fuelled aircraft would produce contrails at lower altitudes due to higher exhaust water content, but that these contrails would likely be less persistent and exhibit lower global warming characteristics due to the absence of soot particles, with potential for complete mitigation through trajectory management [34].

While the foregoing historical developments focused on hydrogen combustion in gas turbine engines, parallel research efforts have explored hydrogen fuel cell propulsion systems as an alternative pathway for aircraft decarbonization. Fuel cells offer the potential for high electrical conversion efficiency, near-zero emissions (producing only water vapor), and significantly reduced noise levels compared to conventional propulsion systems, though at the cost of lower specific power and the need for electric motor propulsion chains [119]. These limitations in power density and energy conversion chain efficiency restrict fuel cell applications primarily to short-range regional and general aviation aircraft [67]. Notable European achievements include the DLR Antares DLR-H2, the first piloted aircraft powered solely by fuel cells [62], and the HY4, a four-seat fuel cell aircraft developed through collaboration between DLR, H2FLY, and other partners [69]. Major European research initiatives have included the Clean Sky Joint Undertaking's investigations into

fuel cell auxiliary power units and hybrid-electric propulsion architectures [129], and dedicated fuel cell aviation projects such as ENFICA-FC, which developed a fuel cell propulsion system for light aircraft [128]. While these developments have demonstrated the technical feasibility of fuel cell propulsion for certain aircraft categories, comprehensive assessments indicate that hydrogen combustion in gas turbines remains the primary pathway for larger, long-range commercial aircraft that demand higher specific power outputs [12]. A comprehensive summary of these efforts can be found in Ref. [149].

1.4 Advantages of Hydrogen Combustion

Beyond the fundamental goal of eliminating direct CO_2 emissions, hydrogen combustion offers several additional advantages that enhance its appeal for aviation applications.

The high gravimetric energy density of hydrogen, as noted previously, means that less fuel mass is required for a given mission. This mass reduction can propagate through the aircraft design, potentially allowing for lighter structures, smaller landing gear, and reduced structural loads. These secondary mass savings can partially offset the penalties associated with the bulky fuel storage system.

Hydrogen's wide flammability range not only enables lean combustion for NO_x reduction but also improves engine reliability. The ability to sustain combustion at very lean conditions enhances altitude relight capability, an important safety consideration for high-altitude cruise. Furthermore, hydrogen's rapid ignition characteristics reduce the risk of flame extinction during transient maneuvers or turbulent flow conditions.

The cryogenic nature of liquid hydrogen opens unique opportunities for thermal management within the propulsion system. The substantial heat sink capacity of LH_2 can be exploited to precool or intercool compressor air, cool turbine cooling air, or heat the fuel itself before combustion, thereby modifying the thermodynamic cycle in ways that can improve overall efficiency. Concepts such as regenerative fuel heating, in which hydrogen is warmed by hot turbine exhaust gases before entering the combustor, and expander cycles, in which hydrogen drives a dedicated turbine to provide auxiliary power, have been shown to offer modest performance improve-

ments compared to conventional Brayton cycles [23]. These alternative cycles are explored in detail in Chapter 2, Section 2.5, where integrated propulsion system modelling demonstrates their potential benefits for hydrogen turbofan engines.

From a broader energy system perspective, hydrogen offers the possibility of decoupling aviation energy supply from fossil fuels. If produced via renewable electricity, hydrogen enables aviation to tap into the growing renewable energy infrastructure, providing a pathway for the sector to participate in the transition to sustainable energy systems. This decoupling also reduces the aviation industry's exposure to the volatility of fossil fuel prices and geopolitical risks associated with oil supply.

1.5 Challenges and Safety Concerns

Despite its advantages, hydrogen combustion introduces a set of technical challenges that must be addressed to enable safe and reliable aviation operations. These challenges stem from hydrogen's fundamental physical and chemical properties and have profound implications for combustor design, engine integration, and aircraft configuration.

1.5.1 Flashback Susceptibility

Flashback is the phenomenon wherein a flame propagates upstream from its intended stabilization location into the fuel-air mixing region. In premixed hydrogen combustors, flashback poses a severe safety hazard because it can lead to flame stabilization in regions not designed to withstand high temperatures, such as fuel injectors, mixing ducts, or even fuel lines. This can cause hardware damage, structural failure, or in the worst case, catastrophic engine failure.

Several physical mechanisms can trigger flashback in hydrogen burners [17]. Boundary layer flashback occurs when the flame propagates upstream through the low-velocity region near the wall, where the flow velocity can fall below the turbulent flame speed. This mechanism is particularly relevant for hydrogen due to its high laminar flame speed and diffusivity. Combustion-induced vortex breakdown (CIVB) involves the disruption of swirling flows due to heat release and density

changes associated with combustion, creating low-velocity regions where flames can anchor and propagate upstream. Autoignition-induced flashback can occur at elevated temperatures and pressures typical of gas turbine combustors, where the short ignition delay time of hydrogen can lead to spontaneous ignition in the premixing region before the intended flame stabilization zone.

Assessing flashback susceptibility requires understanding the balance between local flow velocities and flame propagation speeds. For the wall-flashback, as an example, the critical velocity gradient, defined as the gradient of axial velocity in the radial direction at the wall, is often used as a flashback criterion: flashback is avoided when the velocity gradient exceeds a critical threshold determined by the fuel properties and flow conditions [17]. However, for complex three-dimensional swirling flows typical of practical combustors, simple gradient-based criteria may be insufficient, necessitating more sophisticated analysis techniques such as flame anchoring analysis or direct simulation of flame dynamics.

Prevention of flashback is a central design consideration for hydrogen combustors. Strategies include employing converging nozzles to accelerate the flow and increase velocity gradients near the wall, staging the fuel injection to avoid premixed conditions in vulnerable regions, and carefully controlling swirl intensity and distribution to avoid vortex breakdown. The micro-mixing concept, in which fuel is injected through a large number of small orifices into high-velocity air streams to achieve rapid mixing with minimal residence time in the premixed state, has shown promise for simultaneously reducing NO_x emissions and mitigating flashback risk [45]. This concept is investigated in Chapter 3, Section 3.2, where the effects of converging nozzle geometry on flashback prevention and pollutant formation are explored.

1.5.2 Nitrogen Oxides Formation

While hydrogen combustion eliminates CO_2 emissions, it does not eliminate all environmental impacts. On top of contrail formation, nitrogen oxides, comprising primarily nitric oxide (NO) and nitrogen dioxide (NO_2), are the principal pollutants produced by hydrogen-air combustion. NO_x emissions contribute to ground-level ozone formation, smog, acid rain, and respiratory health problems, and they are subject to increasingly stringent regulatory limits.

The formation of NO_x in hydrogen flames is dominated by the thermal mechanism, also known as the Zeldovich mechanism, in which molecular nitrogen from the air reacts with oxygen atoms at high temperatures to form NO. The rate of thermal NO_x formation increases exponentially with temperature, making it highly sensitive to peak flame temperatures. Since hydrogen-air flames can reach higher adiabatic temperatures than kerosene-air flames, as shown in Figure 1.1, hydrogen combustion has the potential to produce significant NO_x emissions if not carefully controlled.

Fortunately, hydrogen's wide flammability range enables operation at very lean equivalence ratios, which substantially reduces peak flame temperatures and thereby suppresses NO_x formation. Lean premixed combustion has emerged as the preferred strategy for low-NO_x hydrogen burners, accepting the increased flashback risk in exchange for dramatically reduced emissions. Achieving stable, efficient combustion at ultra-lean conditions requires careful attention to flame stabilization, mixing uniformity, and combustion dynamics.

Water or steam injection represents an additional strategy for NO_x mitigation in hydrogen combustors. The introduction of water into the combustion process reduces flame temperatures through three mechanisms: a cooling effect from the latent heat of vaporization and sensible heat absorption by the water, a dilution effect that reduces the concentrations of oxygen and nitrogen in the reaction zone, and a chemical effect in which water participates in radical reactions that alter the kinetic pathways of NO formation. Water injection is explored in Chapter 4, Section 4.2.3, where a chemical reactor network methodology is used to accurately predict NO_x emissions under various water injection loads.

The development of accurate and computationally efficient methods for predicting NO_x emissions is a critical need for hydrogen combustor design. Traditional computational fluid dynamics (CFD) approaches require transporting NO_x species through the entire flow field or post-processing frozen flow solutions with detailed chemistry, both of which are computationally expensive and impractical for extensive parametric studies during the design phase. This thesis introduces a novel chemical reactor network (CRN) methodology that enables rapid and reliable NO_x prediction with minimal calibration data, as detailed in Chapter 4.

1.5.3 Thermoacoustic Instabilities

Thermoacoustic instabilities arise from the coupling between unsteady heat release in the combustion process and acoustic waves in the combustor. When the fluctuations in heat release are in phase with pressure oscillations, energy is fed into the acoustic field, leading to self-sustained oscillations that can reach destructive amplitudes. These instabilities are characterised by intense pressure fluctuations, potentially causing mechanical vibrations, structural fatigue, and reduced component lifetimes.

Lean premixed hydrogen flames exhibit complex sensitivity to thermoacoustic instabilities. The high reactivity and short residence times associated with hydrogen combustion can lead to strong coupling between flow disturbances and heat release fluctuations. The wide flammability range of hydrogen means that equivalence ratio fluctuations, which naturally occur in turbulent flows, can produce significant variations in flame speed and heat release rate, further exacerbating the potential for instability.

Mitigating thermoacoustic instabilities requires careful attention to combustor acoustic characteristics, flame dynamics, and fuel injection strategies. Techniques include passive control measures such as acoustic dampers or careful tuning of combustor length to avoid resonant frequencies, and active control strategies that modulate fuel flow or air split ratios in response to measured pressure oscillations. Understanding the onset conditions and mode shapes of thermoacoustic instabilities is essential for robust combustor design, though this topic is beyond the primary scope of this thesis.

1.5.4 Operational and Logistical Challenges

Beyond the technical challenges associated with combustion and propulsion, hydrogen aviation faces significant operational and logistical hurdles. The cryogenic nature of liquid hydrogen requires specialized storage, handling, and refueling infrastructure at airports. LH_2 must be maintained at approximately 20 K, necessitating vacuum-insulated tanks, cryogenic transfer lines, and careful management of boil-off during ground operations. The development of this infrastructure represents a substantial capital investment and requires coordination across the aviation value chain, from fuel production and distribution to airport operations and aircraft maintenance.

Safety protocols for hydrogen handling are more stringent than those for conventional jet fuel. Hydrogen is highly flammable across a wide range of concentrations, and its low ignition energy makes it susceptible to accidental ignition from static discharge or hot surfaces. Hydrogen embrittlement of certain metals is a concern for long-term material integrity, requiring careful selection of materials for fuel system components. Leak detection systems and rigorous maintenance procedures are essential to ensure safe operations.

The volumetric storage penalty of hydrogen has implications for aircraft design and operations. The larger fuel tanks may necessitate alternative aircraft configurations, as discussed in Section 1.2.3, which in turn affect manufacturing, maintenance, and airport compatibility. The potential need for new aircraft designs, rather than retrofitting existing platforms, increases the development cost and timeline for hydrogen aviation adoption.

Certification of hydrogen aircraft and engines by aviation authorities such as the Federal Aviation Administration (FAA) or the European Union Aviation Safety Agency (EASA) will require extensive testing and demonstration of safety and reliability. Existing certification frameworks are based on decades of experience with kerosene-fuelled aircraft, and adapting these frameworks to hydrogen will likely require new standards and test procedures.

Despite these challenges, the aviation industry and regulatory bodies are increasingly recognizing the necessity of exploring hydrogen as a pathway to sustainable aviation, and efforts are underway to address these operational and logistical issues in parallel with the technical research aspects, as the ones presented in this thesis.

1.6 Research Objectives and Thesis Scope

The overarching goal of this thesis is to develop and demonstrate computational methodologies that enable efficient, accurate, and practical analysis of hydrogen combustion in jet engines, with particular emphasis on addressing the challenges outlined in the preceding sections. Specifically, this research aims to:

1. Perform an integrated system-level analysis of hydrogen-fuelled turbofan engines using the Numerical Propulsion System Simulation (NPSS) software, investigating alternative thermodynamic cycles for cryogenic hydrogen uti-

lization and assessing their performance benefits relative to conventional kerosene-fuelled architectures.

2. Investigate hydrogen flame behaviour through computational fluid dynamics simulations, initially applied to a simple micromix injector configuration to understand fundamental combustion characteristics and flashback mitigation strategies, and subsequently extended to the AHEAD combustor, a lean premixed swirl-stabilized burner featuring complex geometric characteristics including radial swirl generation, axial air injection, mixing tube, and dilution orifices.

3. Characterize flame stabilization behaviour in the AHEAD combustor across a range of operating conditions, with flame position serving as a key indicator of flashback propensity, employing multi-fidelity modelling approaches that integrate experimental measurements with computational simulations to enable efficient design space exploration.

4. Develop a CRN methodology employing optimization algorithms for rapid prediction of NO_x emissions without requiring detailed knowledge of the combustor flow field, validated initially on a non-premixed hydrogen combustor with water injection for emissions reduction, and subsequently applied to the premixed AHEAD configuration.

5. Identify and outline the pathway for integrating the CRN methodology with system-level propulsion modelling tools such as NPSS, establishing the conceptual basis for a unified workflow enabling rapid, emissions-aware preliminary design of hydrogen-fuelled aircraft engines.

The scope of this thesis encompasses computational fluid dynamics simulations using Reynolds-Averaged Navier-Stokes (RANS) methods, multi-fidelity modelling combining computational and experimental data sources, zero-dimensional CRN modelling, and system-level thermodynamic cycle analysis. The research progresses systematically from simplified configurations to industrially relevant geometries, establishing modelling capabilities of increasing complexity. While high-fidelity techniques such as Large Eddy Simulation (LES) and experimental testing provide reference data for validation, the primary emphasis is on developing practical, computationally efficient methods that balance accuracy with the need for rapid iteration

during the design process. This research does not address all aspects of hydrogen aviation. Topics such as detailed aircraft aerodynamic design, structural analysis of hydrogen storage systems, hydrogen production and distribution infrastructure, economic viability assessments, and regulatory frameworks are outside the scope of this thesis, though their importance is acknowledged. The focus is squarely on combustion and propulsion system modelling for hydrogen-fuelled aircraft, with the understanding that successful implementation of hydrogen aviation will require coordinated advances across multiple disciplines.

1.7 Thesis Outline

The remainder of this thesis is organized into three main chapters, followed by conclusions and recommendations for future work.

Chapter 2 provides an overview of hydrogen properties relevant to hydrogen combustor design and analysis, such as the laminar flame speed, the adiabatic flame temperature and the specific heat of combustion products. Then, the mechanisms for the generation of nitric oxide during hydrogen combustion with air are presented, showing with a counterflow flame analysis that water could be successfully injected into the combustion chamber to abate emissions. Finally, the chapter describes the NPSS framework for system-level analysis of gas turbine engines, presenting a comparison between a kerosene and a hydrogen-fuelled turbofan, including assessment of alternative thermodynamic cycles and cryogenic fuel system design.

Chapter 3 presents two CFD studies using RANS. The chapter begins with a summary of the governing equations and turbulence-chemistry interaction model. Then, the first work describes the numerical investigation of a simple hydrogen micro-mixing injector, focusing on the effects of introducing a converging nozzle on flashback prevention and pollutant formation. Then, 2D axisymmetric RANS are used to study flame position inside the AHEAD swirl-stabilized lean premixed hydrogen combustor. A multi-fidelity model, developed in Optimad [110] for efficient design space exploration, is then used to improve the predictions obtained by CFD with a limited number of experimental points.

Chapter 4, which constitutes the core contribution of this thesis, presents the CRN methodology for rapid NO_x prediction, using an approximation of the characteristic

flow zones and an optimization algorithm to replicate high-fidelity outputs. The chapter includes two study cases: the first applies the methodology to a non-premixed hydrogen burner, demonstrating its capability to accurately capture the effects of water injection on NO_x production; the second applies the methodology to the AHEAD premixed hydrogen burner, showing its effectiveness in replicating the experimental data.

Chapter 5 concludes the thesis by first summarising the principal results achieved across the three modelling fronts: the accuracy of the CRN methodology for NO_x prediction at negligible computational cost, the performance of the multi-fidelity framework for flashback assessment, and the cross-domain applicability demonstrated through the application to hybrid rocket engines. Two priority directions for future research are then identified. The first addresses the generalisation of the CRN methodology to operating regimes in which the flow topology is documented to change and to configurations requiring turbulence–chemistry interaction modelling beyond the perfectly stirred reactor assumption. The second proposes the deep integration of the CRN into the NPSS engine model, establishing a unified workflow for simultaneous thermodynamic cycle optimization and emissions assessment that enables emissions-conscious preliminary design of hydrogen aviation propulsion systems.

Chapter 2

Hydrogen Combustion and Propulsion System Integration

2.1 Introduction

Chapter 1 introduced hydrogen as a promising aviation fuel, highlighting its high gravimetric energy density, zero direct carbon emissions, and the principal challenges associated with its implementation. This chapter provides a quantitative examination of the combustion properties that are relevant to hydrogen burner design and performance. The discussion builds upon the qualitative overview presented in Section 1.2.1 by establishing relationships between operating conditions and combustion behaviour, supported by equilibrium and flame calculations performed using Cantera [50].

The chapter is organized as follows. Section 2.2 examines the combustion characteristics of hydrogen, including laminar flame speed, transport properties, and their variation with equivalence ratio, pressure, and temperature. Section 2.3 presents the detailed nitrogen oxide formation mechanisms relevant to hydrogen combustion, extending the overview in Section 1.5.2. Finally, Section 2.4 discusses the opportunities that hydrogen presents for aero gas turbine propulsion, studying an ultrahigh bypass ratio turbofan in the context of a blended wing body aircraft [40].

2.2 Combustion Properties of Hydrogen

The combustion behaviour of hydrogen differs fundamentally from that of hydrocarbon fuels due to its molecular structure and reaction kinetics. This section quantifies these differences in the conditions relevant to aviation gas turbine combustors.

2.2.1 Laminar Flame Speed

The laminar flame speed (S_L) represents the velocity at which a planar flame front propagates into an unburned mixture under laminar flow conditions. This fundamental property governs flame stabilization, turbulent flame speed correlations, and flashback susceptibility.

Effect of Equivalence Ratio

Figure 2.1 presents the laminar flame speed as a function of equivalence ratio for hydrogen and a Jet-A surrogate (n-dodecane) at atmospheric pressure and an unburned mixture temperature of 300 K. Calculations were performed using both the CRECK mechanisms for hydrogen [104, 75, 10] and for n-dodecane [121, 140, 141], without considering N_2 oxidation to reduce the computational cost.

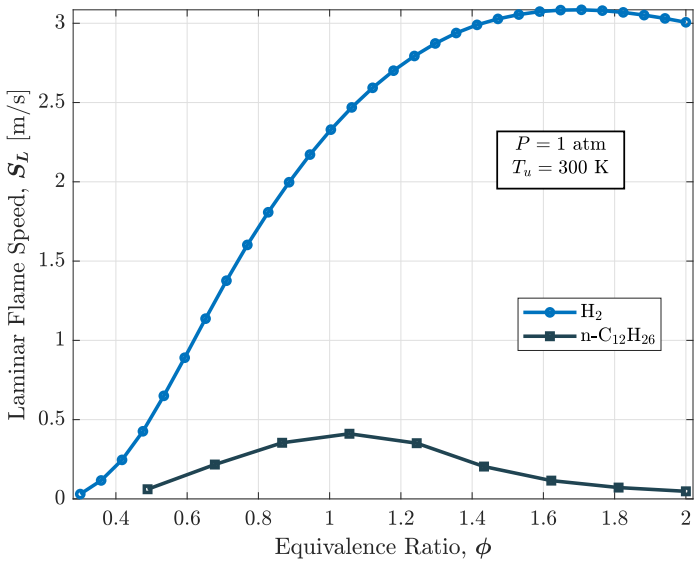

Fig. 2.1 Laminar flame speed as a function of equivalence ratio for hydrogen, methane, and Jet-A surrogate (n-dodecane) at atmospheric pressure and T_u = 300 K.

Several observations emerge from this comparison:

1. The maximum laminar flame speed of hydrogen (approximately 3.1 m/s near $\phi = 1.7$) exceeds that of the jet-A surrogate, always below 0.5 m/s, by nearly an order of magnitude.

2. Even at the lean equivalence ratios relevant to low-NO_x combustion ($\phi = 0.4$–0.6), hydrogen maintains laminar flame speeds of 0.5–1.5 m/s, still substantially higher than the peak values for hydrocarbons.

3. The extended lean flammability limit of hydrogen (down to $\phi \approx 0.1$) compared to methane ($\phi \approx 0.5$) enables ultra-lean operation strategies discussed in Section 1.5.2.

Effect of Pressure

Figure 2.2 shows the laminar flame speed of hydrogen-air mixtures at pressures from 1 to 50 bar, spanning the range relevant to aircraft engine combustors from idle to maximum thrust conditions and exhibiting a reduction in laminar flame speed against pressure in this range.

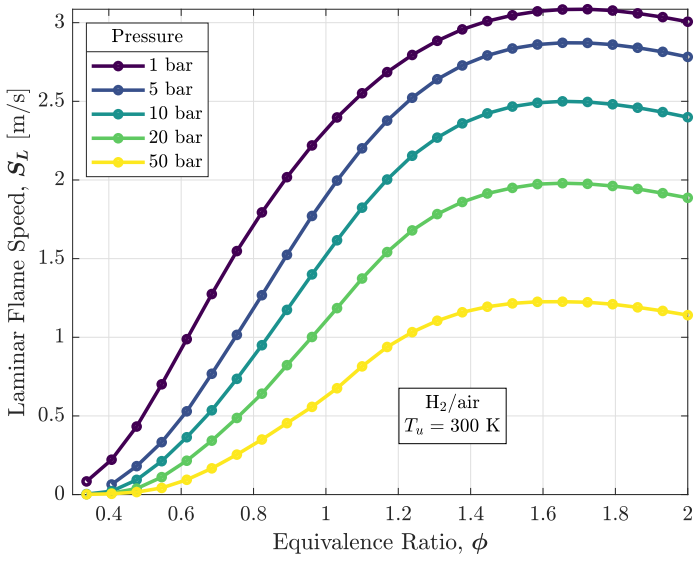

Fig. 2.2 Laminar flame speed of hydrogen-air mixtures as a function of equivalence ratio at different pressures (1–50 bar) and T_u = 300 K.

Effect of Unburned Gas Temperature

The temperature of the unburned mixture entering the flame front substantially affects the laminar flame speed. In gas turbine combustors, the air temperature at the combustor inlet varies from approximately 400 K at idle to over 800 K at maximum thrust, depending on the overall pressure ratio. Figure 2.3 illustrates this effect for stoichiometric and lean hydrogen-air mixtures.

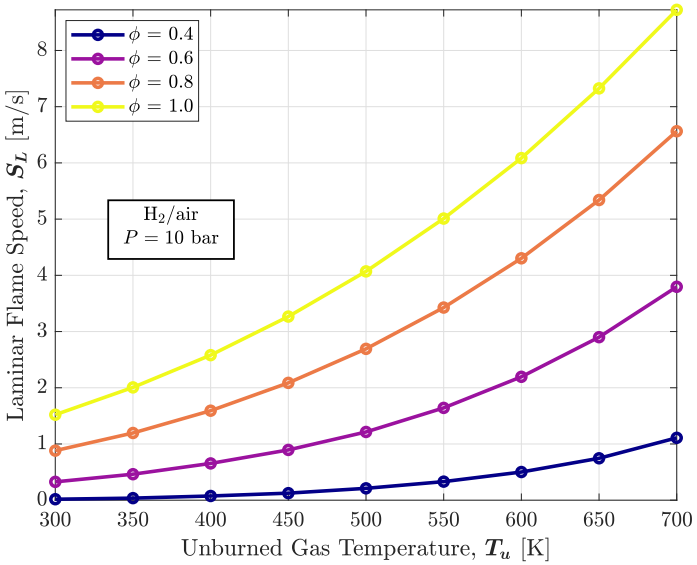

Fig. 2.3 Laminar flame speed of hydrogen-air mixtures as a function of unburned gas temperature at 10 bar for various equivalence ratios. Higher inlet temperatures significantly increase flame speed, affecting flashback margins.

The strong increase in flame speed with temperature has important operational implications. At high-power conditions where both pressure and temperature are elevated, the laminar flame speed can increase by a factor of 2–3 compared to sea-level static conditions, requiring careful attention to flashback margins across the operating envelope. This sensitivity to unburned gas temperature is especially significant for hydrogen-powered aircraft. As discussed in Section 2.4.1, the adoption of advanced thermodynamic cycles, such as regenerative fuel heating, results in gaseous hydrogen entering the combustor at significantly elevated temperatures.

Diffusive Stability and Flame Structure

Hydrogen's molecular transport characteristics differ drastically from those of conventional hydrocarbon fuels. The binary diffusion coefficient of hydrogen in air exceeds that of typical Jet-A vapour surrogates by approximately an order of magnitude across relevant temperature ranges [98, 118]. This elevated diffusivity, combined with hydrogen's low molecular weight, yields a Lewis number, i.e. the ratio between thermal and mass diffusivity, substantially below unity. In contrast, kerosene vapour exhibits Lewis numbers well above unity due to its heavier molecular constituents and correspondingly lower diffusion rates.

The low Lewis number regime introduces thermo-diffusive instabilities in premixed hydrogen flames. Preferential diffusion of the deficient reactant into high-curvature flame regions locally enriches the mixture, elevating reaction rates and producing super-adiabatic flame temperatures in regions of positive curvature [7]. This mechanism fundamentally alters flame morphology and stabilization dynamics compared to the thermo-diffusively stable regime characteristic of kerosene combustion. The practical consequence is increased susceptibility to cellular flame structures and altered sensitivity to stretch effects, both of which influence flame anchoring and flashback propensity [81].

Ignition Behaviour

Ignition characteristics present a paradoxical safety profile that distinguishes between static and dynamic operational scenarios. From a static perspective, hydrogen appears inherently safer: its autoignition temperature (AIT) of approximately 858 K substantially exceeds that of kerosene at approximately 483 K [65]. This suggests reduced susceptibility to unintended ignition in quiescent, thermally uniform environments.

However, dynamic ignition behaviour reveals a contrasting vulnerability. Hydrogen's minimum ignition energy (MIE) of approximately 0.017 mJ falls an order of magnitude below that of large hydrocarbons [154]. This extreme sensitivity to electrostatic discharge, mechanical sparks, and other transient energy sources introduces operational hazards in fuel handling, storage, and distribution systems. The apparent contradiction between high AIT and low MIE reflects the distinction between thermally driven ignition mechanisms (governed by chain-branching kinetics at elevated

temperatures) and spark-induced ignition (governed by radical initiation in localized high-energy zones).

Furthermore, the temperature regime between approximately 700 K and 900 K, typical of compressor discharge conditions in modern gas turbines, reveals a fundamental kinetic difference. Kerosene exhibits negative temperature coefficient (NTC) behaviour in this range, where reaction rates locally decrease with increasing temperature due to competition between chain-branching and chain-terminating pathways [26]. This NTC behaviour increases kerosene's reactivity relative to hydrogen within this window. However, at higher temperatures, such as those encountered near hot surfaces in a combustor, hydrogen's ignition delay decreases drastically. This suggests that in the premixing section of a hydrogen burner, stringent control of potential hotspots is required to prevent ignition and subsequent flashback.

Specific Heat of Combustion Products

Figure 2.4 presents the specific heat at constant pressure (c_p) of the equilibrium combustion products as a function of temperature for hydrogen-air and Jet-A-air combustion at different equivalence ratios.

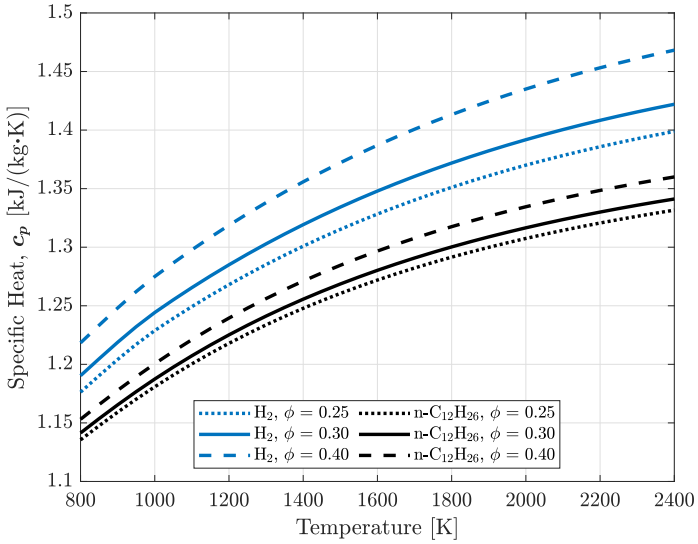

Fig. 2.4 Specific heat at constant pressure of equilibrium combustion products as a function of temperature for stoichiometric hydrogen-air and Jet-A-air mixtures at 10 bar. The higher water content in hydrogen products results in elevated c_p.

The elevated c_p of hydrogen combustion products (approximately 5–10% higher than kerosene products at turbine-relevant temperatures) has direct implications for turbine performance. For a given temperature drop across the turbine:

$$w_t = c_p \Delta T_t \tag{2.1}$$

the specific work output is proportionally higher. This effect contributes to the thermodynamic advantages of hydrogen discussed in Section 2.4.

Adiabatic Flame Temperature at Elevated Pressures

While Figure 1.1 in Chapter 1 illustrated the adiabatic flame temperature at a fixed reference pressure, actual combustor operating conditions vary significantly across the flight envelope. Figure 2.5 extends this analysis by considering the specific air temperature and pressure at the burner inlet for two key operating points: cruise and take-off.

Assuming an Overall Pressure Ratio (OPR) of 50 and isentropic compression, the burner inlet conditions for cruise (35,000 ft) are approximately 670 K and 12 bar. For sea-level take-off, these conditions increase to approximately 880 K and 51 bar. To ensure a consistent comparison between hydrogen and the Jet-A surrogate, the fuel injection temperature is kept constant at 300 K for all cases.

The results highlight that the potential for NO_x production is significantly higher during take-off. This is driven by the combination of elevated burner inlet temperatures and the higher equivalence ratios typically required during high-thrust operations.

2.3 Nitrogen Oxide Formation Mechanisms

Section 1.5.2 introduced the thermal NO_x formation pathway and the strategy of lean combustion for emissions reduction. This section provides a more comprehensive treatment of the nitrogen chemistry relevant to hydrogen combustion.

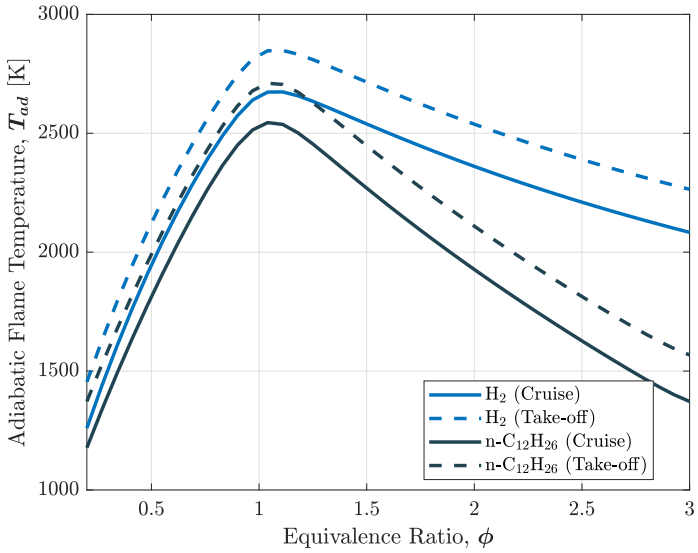

Fig. 2.5 Adiabatic flame temperature of hydrogen-air mixtures as a function of equivalence ratio for cruise and take-off burner inlet conditions.

2.3.1 Thermal NO (Extended Zeldovich Mechanism)

The thermal or Zeldovich mechanism [161, 162] is the dominant NO_x formation pathway at high temperatures. The mechanism consists of three elementary reactions:

$$O + N_2 \rightleftharpoons N + NO \tag{2.2}$$

$$N + O_2 \rightleftharpoons O + NO \tag{2.3}$$

$$N + OH \rightleftharpoons H + NO \tag{2.4}$$

The first reaction (Eq. 2.2) has a high activation energy, i.e. approximately 315 kJ/mol [14], requiring high temperatures to initiate the chain of reactions [132]. Thus, at temperatures below approximately 1800 K, NO formation through this mechanism is significantly reduced [93].

2.3.2 N_2O Intermediate Pathway

At lower temperatures and in fuel-lean conditions, where the Zeldovich mechanism is not relevant, the N_2O intermediate mechanism can contribute significantly to NO

formation [93]. The principal reactions are:

$$N_2 + O + M \rightleftharpoons N_2O + M \tag{2.5}$$

$$N_2O + O \rightleftharpoons NO + NO \tag{2.6}$$

$$N_2O + O \rightleftharpoons N_2 + O_2 \tag{2.7}$$

$$N_2O + H \rightleftharpoons N_2 + OH \tag{2.8}$$

The three-body nature of the reaction (Eq. 2.5) favours higher pressures [132], making this pathway increasingly important with the elevated pressures characteristic of modern gas turbine combustors. The N_2O pathway is particularly relevant for lean premixed hydrogen combustion, where peak temperatures are suppressed below the threshold for significant thermal NO but pressures remain high.

2.3.3 NNH Pathway

The NNH radical pathway has been identified as potentially significant in hydrogen flames, particularly under fuel-rich or high-hydrogen conditions [19]:

$$N_2 + H \rightleftharpoons NNH \tag{2.9}$$

$$NNH + O \rightarrow NO + NH \tag{2.10}$$

Its importance in hydrogen flames stems from the high concentrations of H atoms in the reaction zone, which are substantially greater than in hydrocarbon flames due to hydrogen's high diffusivity and the radical pool composition [120].

The NNH pathway can be relevant particularly at moderate temperatures where thermal NO is suppressed but H-atom concentrations remain elevated: conditions that could be met in a Rich-Quench-Lean (RQL) hydrogen combustor.

2.3.4 Water and Steam Injection for NO_x Abatement

Beyond lean premixed combustion, water and steam injection represent alternative strategies for NO_x reduction in hydrogen combustors. This approach has historical precedent in gas turbine applications for power augmentation and efficiency enhancement [160], and has recently attracted renewed interest as an emissions

mitigation technique within frameworks such as MILD combustion [25] and exhaust gas recirculation. Contemporary applications include the WET engine configuration proposed by MTU for aerospace applications [68] and the Hydrogen Steam and Inter-Cooled Turbine Engine (HySITE) concept developed by Pratt & Whitney [3].

Mechanisms of NO_x Reduction

The introduction of water into combustion systems affects NO_x formation through three distinct mechanisms [46, 60, 61, 28, 29]:

1. **Thermal effect**: Water absorbs heat through both the latent heat of vaporization (for liquid injection) and the sensible heat capacity of steam. The specific heat of water vapour ($c_p \approx 2.0$ kJ/(kg·K)) exceeds that of air and other combustion products, reducing peak flame temperatures and thereby suppressing thermal NO formation.
2. **Dilution effect**: The addition of steam reduces the partial pressures and mass fractions of reactants, lowering the volumetric heat release rate. This dilution decreases the concentrations of oxygen and nitrogen in the reaction zone, directly reducing the collision frequency of species involved in NO formation.
3. **Chemical effect**: Water participates actively in the reaction kinetics through its high collision efficiency in third-body reactions. As a major combustion product, water also shifts chemical equilibrium conditions, affecting radical concentrations and reaction pathways.

The relative importance of these mechanisms depends on the injection method. For liquid water injection, the thermal effect dominates due to evaporative cooling, while the chemical effect remains secondary. For steam injection, where no phase change occurs, the dilution and chemical effects assume greater importance relative to thermal effects.

Third-Body Reaction Enhancement

A distinctive aspect of water's influence on hydrogen combustion kinetics arises from its high third-body collision efficiency. Early investigations by Koroll and

Mulpuru [78] demonstrated that when hydrogen-oxygen premixed flames are diluted with steam, the reduction in laminar flame speed is less pronounced than would be predicted if water behaved as an inert diluent. This observation indicates that water actively participates in the reaction kinetics beyond simple thermal and dilution effects.

The primary mechanism involves the third-body recombination reaction:

$$H + O_2 + M \rightleftharpoons HO_2 + M \tag{2.11}$$

where water exhibits enhanced collision efficiency compared to nitrogen or other diluents. This reaction competes with the chain-branching reaction H + $O_2 \rightarrow$ OH + O, and the balance between these pathways governs flame propagation characteristics. Studies by Lamioni et al. [80] and Attili et al. [8] have identified this competition as temperature-dependent and sensitive to the presence of species with high third-body efficiencies.

The enhanced third-body efficiency of water also influences NO_x chemistry directly. Investigations by Göke et al. [48, 49] demonstrated that water affects the redistribution of radical species, particularly atomic oxygen (O), which modulates the N_2O intermediate pathway described in Section 2.3.2. Concetti et al. [30] further showed that the chemical effect of water can reduce NO_x concentrations even at constant temperature, indicating that the mechanism extends beyond simple thermal suppression.

Counterflow Flame Analysis

The effects of steam injection on hydrogen flame structure and NO_x formation can be examined through canonical counterflow diffusion flame configurations [152], computed in Cantera [50]. While idealized compared to practical combustors, counterflow flames provide a tractable framework for isolating the fundamental physics of steam-flame interactions.

Figure 2.6 presents the maximum temperature of hydrogen-air counterflow diffusion flames as a function of strain rate for various steam mass fractions. Calculations were performed using the CRECK mechanism [136, 142, 143], which captures both high-temperature hydrogen oxidation and nitrogen chemistry. At low strain rates, the flame approaches chemical equilibrium with temperatures near 2300 K. As strain

rate increases, enhanced mixing reduces residence time in the reaction zone, progressively lowering the maximum temperature until flame extinction occurs. The addition of steam shifts both the temperature profile and the extinction strain rate towards lower values.

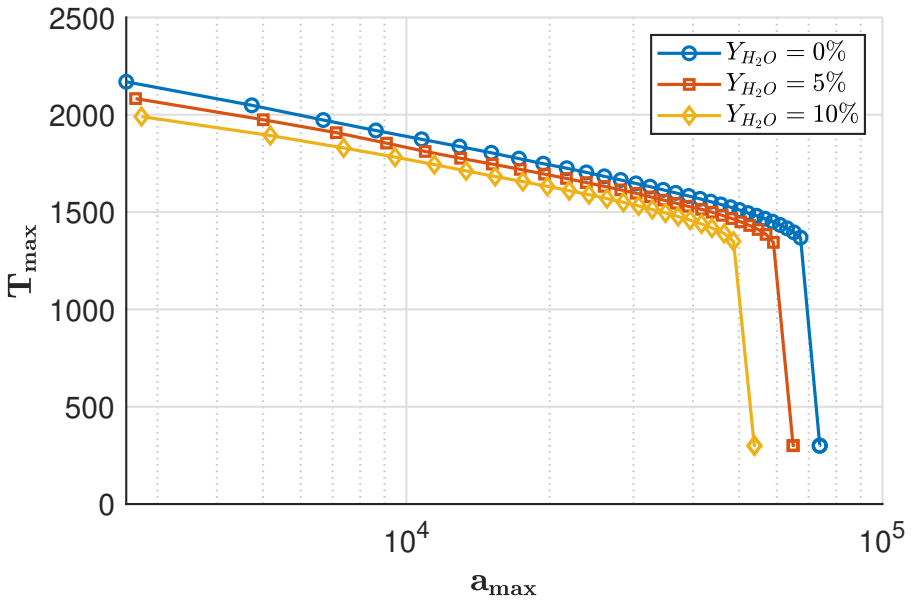

Fig. 2.6 Maximum temperature of hydrogen-air counterflow diffusion flames as a function of strain rate for various steam mass fractions. The rapid temperature decrease at high strain rates indicates flame quenching.

The corresponding NO molar fractions are shown in Figure 2.7. Given the exponential temperature dependence of thermal NO formation, the temperature reductions induced by steam addition translate to substantial NO abatement. The effect is most pronounced at low strain rates, where residence times are sufficient for significant NO accumulation. At high strain rates, NO concentrations remain low regardless of steam content due to kinetically limited residence times.

These results show that steam injection reduces NO_x emissions in non-premixed flame configurations while maintaining flame temperatures within ranges acceptable for gas turbine operation. However, the counterflow configuration represents an idealized canonical flame that does not capture the complex flow structures, mixing patterns, and residence time distributions characteristic of practical combustors. The quantitative validity of counterflow-based predictions for realistic burner geometries remains uncertain, motivating the development of more sophisticated modelling approaches. Chapter 4 addresses this limitation through a chemical reactor network

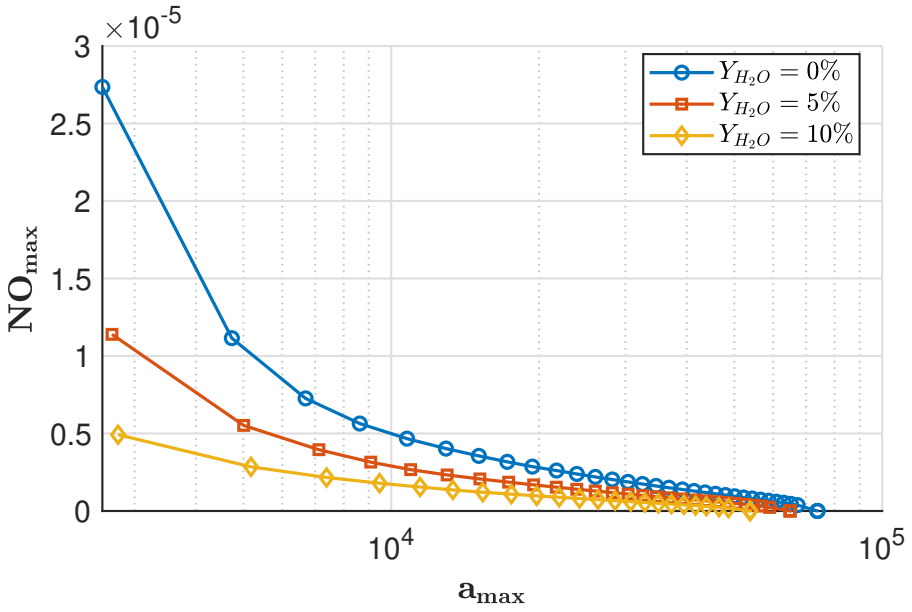

Fig. 2.7 Maximum NO molar fraction in hydrogen-air counterflow diffusion flames as a function of strain rate for various steam mass fractions.

methodology that bridges fundamental kinetic predictions with combustor-level emissions assessment, enabling accurate NO_x prediction under water injection conditions without requiring computationally prohibitive three-dimensional reactive simulations.

2.4 Hydrogen Opportunities for Aero Gas Turbine Propulsion

While the preceding sections have focused on fundamental properties of hydrogen related to combustion, hydrogen also offers unique opportunities that can enhance gas turbine performance and reduce environmental impact beyond the elimination of CO_2 emissions. This section examines these opportunities for an ultra-high bypass turbofan, establishing the context for possible advantages of hydrogen propulsion.

2.4.1 Thermodynamic Cycle Opportunities

The cryogenic storage temperature of liquid hydrogen (approximately 20 K at atmospheric pressure) enables thermodynamic cycle modifications unavailable with ambient-temperature fuels. The heat sink capacity of LH_2 transitioning from storage temperature to combustion temperatures exceeds 4 MJ/kg, substantially greater than the sensible heating of kerosene [23]. This thermal capacity can be exploited to modify the traditional gas turbine Brayton cycle, potentially improving overall propulsion system efficiency.

Alternative Cycle Architectures

Brewer [23] systematically evaluated five alternative cycle concepts that exploit hydrogen's cryogenic properties: compressor air precooling, compressor air intercooling, cooling of turbine cooling air, regenerative fuel heating, and expander cycle. Each approach involves heat exchange between the cryogenic fuel and engine airflow or auxiliary systems, with distinct trade-offs in performance, complexity, and integration challenges.

Compressor air precooling and compressor air intercooling reduce the work required for compression by cooling the air before or between compressor stages. While this approach can improve cycle efficiency, the introduction of heat exchangers in the main airflow path before the combustor results in total pressure losses that can offset the compression work savings. Additionally, heat transfer effectiveness is lower at the relatively low temperatures characteristic of the compressor section.

Turbine cooling air conditioning exploits the cryogenic fuel to precool the air extracted for turbine blade cooling, enabling higher turbine inlet temperatures for a given blade material capability. This approach can increase specific thrust and thermal efficiency, though it requires careful integration with the engine's secondary air system and introduces additional heat exchanger complexity.

Regenerative fuel heating warms the hydrogen fuel before combustion using heat from the engine exhaust gases downstream of the turbine. This cycle modification recovers waste heat that would otherwise be lost through the nozzle, improving overall thermal efficiency by reducing the enthalpy deficit associated with introducing cryogenic fuel directly into the combustor. The fuel preheating reduces the quantity

of energy that must be supplied by combustion to achieve the desired turbine inlet temperature. However, increasing the fuel inlet temperature may facilitate flashback in premixed systems.

Expander cycle extends the regenerative fuel heating concept by routing the heated hydrogen through a dedicated expansion turbine before injection into the combustor. This turbine generates auxiliary power for aircraft systems, eliminating or reducing the need to extract mechanical power from the engine's main shaft through the accessory gearbox. The fuel must be pressurized to higher levels to maintain adequate injection pressure after expansion, requiring more capable fuel pumps but offering the benefit of decoupled auxiliary power generation.

Based on the analysis in Ref. [23], regenerative fuel heating and expander cycle emerge as the most advantageous configurations. Both concepts locate the primary heat exchanger downstream of the low-pressure turbine, where exhaust gas temperatures remain high (facilitating effective heat transfer) while avoiding pressure losses in the combustor inlet flow. The following subsections quantify the performance characteristics of these two cycle architectures through integrated propulsion system modelling.

2.4.2 System-Level Performance Analysis Framework

To quantify the system-level benefits enabled by hydrogen's cryogenic properties and alternative thermodynamic cycles, integrated propulsion system modelling is required. The NPSS software provides a suitable framework for this analysis. NPSS is an object-oriented simulation environment originally developed by NASA Glenn Research Center and now maintained by the Southwest Research Institute [138, 87]. The software connects user-defined components (compressors, turbines, combustors, nozzles, ducts, heat exchangers) through logical links representing mass, momentum, and energy transfer, with each component governed by thermodynamic and mechanical relationships.

For the present analysis, NPSS incorporates thermodynamic property packages enabling precise evaluation of fluid properties under diverse conditions. The CEA package [51] is employed for hydrogen-air combustion calculations, providing chemical equilibrium compositions and thermodynamic properties across the temperature and pressure ranges encountered in gas turbine engines. For cryogenic liquid hy-

drogen properties, the CEA database is supplemented with tabulated data from the CoolProp library [16], which extends property calculations to temperatures below 200 K where CEA data are unavailable.

The solver adjusts independent variables iteratively to satisfy specified requirements. For the two-spool engines analysed here, independent variables include fuel mass flow rate (to achieve the desired turbine inlet temperature), high-pressure turbine pressure ratio, and low-pressure turbine pressure ratio (to maintain energy balance across both spools). The inlet mass flow rate is adjusted to achieve the required thrust level for the specified mission. Only steady-state conditions are considered in this analysis; transient behaviour and dynamic response are beyond the present scope.

For hydrogen engines, additional components model the cryogenic fuel system: a liquid hydrogen tank with thermodynamic state calculations, a cryogenic pump to pressurize the fuel, and heat exchangers to condition the fuel temperature before combustion. The pump efficiency is estimated from empirical correlations in Ref. [107] as a function of specific speed:

$$N_s = \frac{N\sqrt{Q_\mathrm{p}}}{H_\mathrm{stage}^{3/4}} \tag{2.12}$$

where N is the pump rotational speed (rpm), Q_p is the volumetric flow rate (gpm), and H_stage is the head per stage (ft). Heat exchanger effectiveness is specified at $\varepsilon_\mathrm{HX} = 0.8$ following the recommendation in Ref. [23], with a 4% total pressure drop across the air side of the heat exchanger.

This modelling framework enables direct comparison between kerosene and hydrogen propulsion systems operating under identical mission requirements, isolating the effects of fuel properties and cycle architecture on overall performance.

2.5 Propulsion System Performance with Hydrogen

This section presents a systematic performance assessment of hydrogen-fuelled turbofan engines employing the alternative thermodynamic cycles described in Section 2.4.1. The analysis compares baseline kerosene propulsion against three

hydrogen configurations: an idealized case with preheated fuel, regenerative fuel heating, and expander cycle. Mission requirements are derived from a representative long-range BWB aircraft configuration investigated in [71].

2.5.1 Reference Aircraft and Mission Requirements

The design of hydrogen-fuelled aircraft is driven by the need to minimize drag penalties associated with the large fuel storage volumes required for liquid hydrogen. BWB architectures offer substantial internal volume while maintaining aerodynamic efficiency, making them attractive candidates for hydrogen propulsion. The reference configuration is taken from Karpuk et al. [71], which presents the preliminary design of a long-range BWB transport aircraft with three ultra-high bypass ratio (UHBPR) turbofan engines.

The thrust requirement and engine specifications align with the "Large Turbofan" configuration of the European Union-funded ENOVAL project [103]. Cruise conditions are Mach 0.79 at 7650 m altitude. The required thrust per engine at mid-cruise is calculated from:

$$F_{\text{mid cruise}} = \frac{m_{\text{mid cruise}} \cdot g}{N_{\text{engines}} \cdot E_{\text{LD,cruise}}} \tag{2.13}$$

where $m_{\text{mid cruise}}$ is the aircraft mass during mid-cruise (estimated by subtracting half the fuel weight from maximum takeoff weight), $N_{\text{engines}} = 3$ is the number of engines, and $E_{\text{LD,cruise}}$ is the lift-to-drag ratio during cruise. This yields a required thrust of 39.76 kN (8938.6 lbf) per engine at cruise conditions.

The engine architecture is a two-spool, separate-exhaust turbofan with bypass ratio (BPR) of 16.2 at cruise, consistent with UHBPR technology. Component efficiencies are specified in Table 2.1 based on Ref. [23], with combustor efficiency assumed to be 0.98. The turbine inlet temperature (TIT) is fixed at 1600 K, representative of modern commercial turbofan engines at cruise [32].

While the ENOVAL specifications provide BPR at cruise, the values of overall pressure ratio (OPR) and fan pressure ratio (FPR) are given only at top-of-climb conditions in Ref. [103]. To determine appropriate design parameters for cruise operation, a parametric sensitivity analysis is performed to identify combinations of OPR and FPR that minimize thrust-specific fuel consumption (TSFC) while

Table 2.1 Engine component efficiencies for propulsion system analysis, derived from [23].

Component	Efficiency
Fan adiabatic efficiency	0.892
Low-pressure compressor adiabatic efficiency	0.865
High-pressure compressor adiabatic efficiency	0.862
Combustor efficiency	0.98
Combustor pressure drop	0.045
High-pressure turbine adiabatic efficiency	0.90
Low-pressure turbine adiabatic efficiency	0.88

satisfying the thrust requirement at fixed BPR and TIT. This approach follows standard gas turbine cycle analysis methodology [99].

A schematic diagram of the baseline two-spool separate-exhaust turbofan configuration is shown in Figure 2.8, with the numbering convention for engine stations.

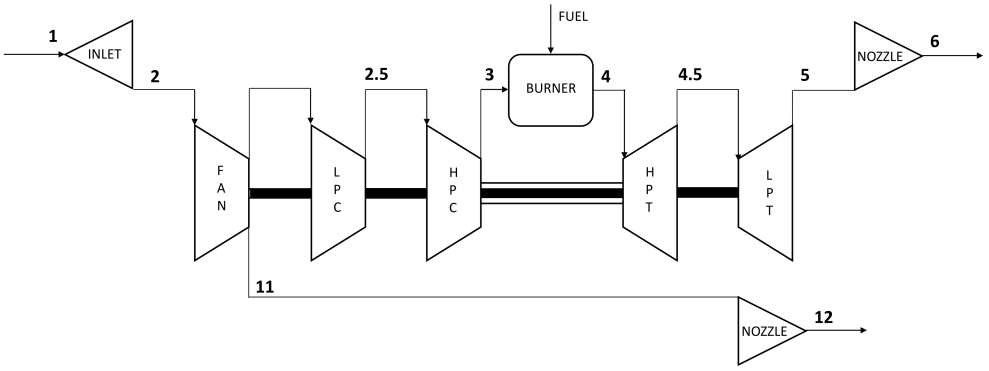

Fig. 2.8 Schematic diagram of the two-spool separate-exhaust turbofan engine configuration analysed in this study.

2.5.2 Baseline Kerosene Engine Performance

The parametric analysis for the kerosene-fuelled engine explores the TSFC variation with specific thrust as functions of OPR and BPR for fixed FPR and TIT. Representative results are shown in Figure 2.9 for four selected FPR values. Each curve corresponds to a constant OPR, with BPR varying along the curve. The analysis reveals that for each OPR, an optimal FPR exists that minimizes TSFC.

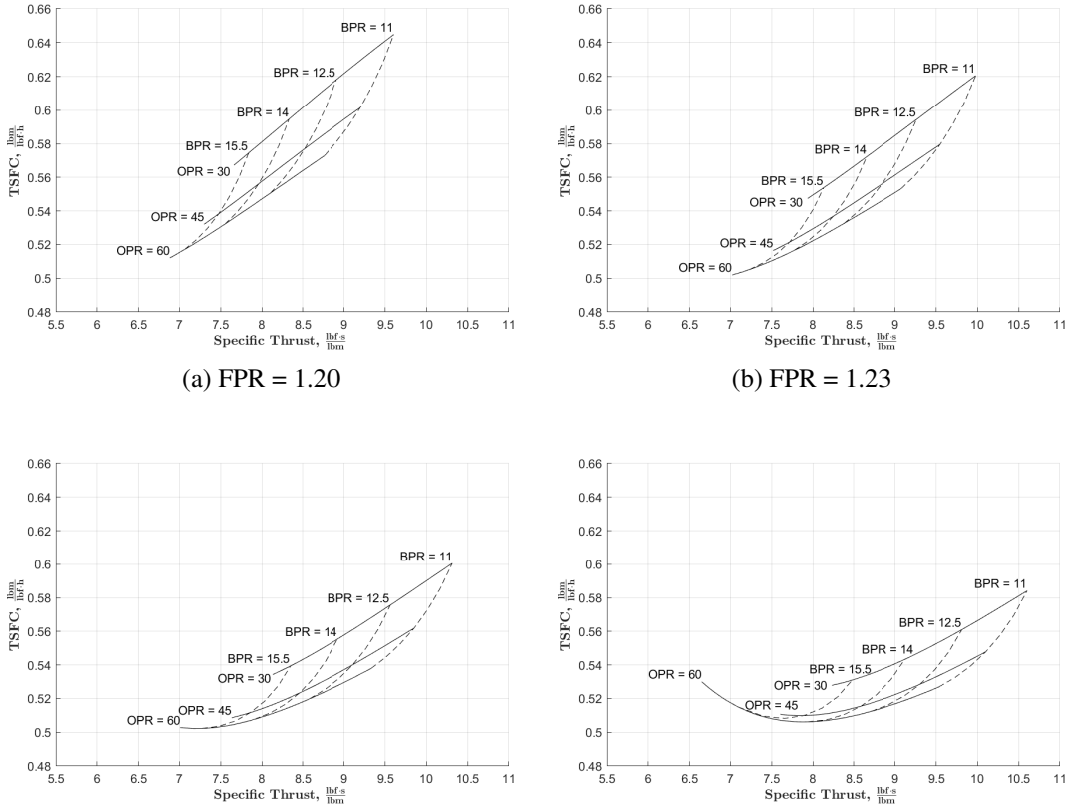

(a) FPR = 1.20

(b) FPR = 1.23

(c) FPR = 1.26

(d) FPR = 1.29

Fig. 2.9 Parametric analysis of kerosene-fuelled turbofan showing TSFC versus specific thrust for varying OPR and BPR at fixed TIT = 1600 K and selected FPR values.

With BPR fixed at 16.2, a two-dimensional optimization over OPR and FPR identifies the minimum TSFC configuration. Figure 2.10 presents the TSFC surface as a function of these two design parameters. The optimal point occurs at FPR = 1.23 and OPR = 67.5, yielding TSFC = 0.4995 lbm/(lbf·h) with a corresponding kerosene mass flow rate of 0.562 kg/s to achieve the required thrust at the specified TIT.

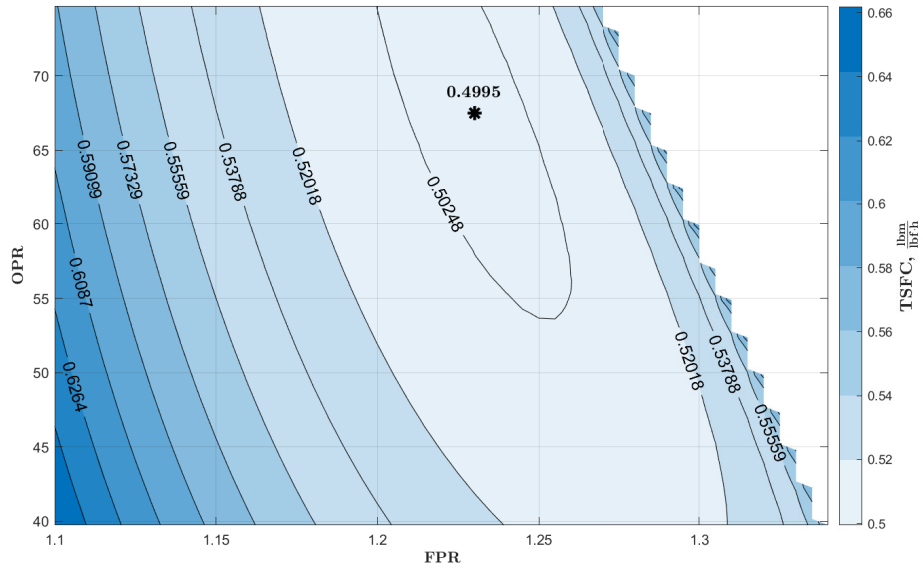

Fig. 2.10 Kerosene-fuelled engine: TSFC as a function of OPR and FPR with fixed BPR = 16.2 and TIT = 1600 K. The optimal configuration minimizing TSFC is indicated.

This baseline performance establishes the reference against which hydrogen configurations are compared. However, direct comparison of TSFC between engines operating on different fuels can be misleading due to the substantial difference in fuel heating values. A more appropriate metric is the thrust-specific energy consumption (TSEC), defined as:

$$\text{TSEC} = \text{TSFC} \times \text{LHV} = \frac{\dot{m}_{\text{fuel}} \cdot \text{LHV}}{F} \tag{2.14}$$

where LHV is the lower heating value of the fuel. For kerosene, LHV $\approx$ 43 MJ/kg; for hydrogen, LHV $\approx$ 120 MJ/kg [2]. TSEC represents the energy consumption per unit thrust, enabling meaningful fuel-agnostic performance comparisons.

2.5.3 Hydrogen Engine with Preheated Fuel

Before analyzing the integrated cycles with heat exchangers, an idealized hydrogen configuration is examined in which the fuel enters the combustor at ambient temperature (298 K) after being heated by an unspecified external heat source. This represents an upper bound on performance, isolating the benefits of hydrogen's combustion properties from the penalties associated with cryogenic fuel introduction and heat exchanger integration.

The parametric analysis procedure is repeated for this hydrogen configuration, maintaining the same component efficiencies, BPR, and TIT as the kerosene baseline. Figure 2.11 presents the TSFC surface. The optimal design point shifts to FPR = 1.25 and OPR = 71.3, reflecting the altered thermodynamic properties of hydrogen combustion products. As discussed in Section 2.2.1, the higher water content in hydrogen exhaust gases results in elevated specific heat, which increases the specific work extractable from the turbine for a given temperature drop. This effect modifies the optimal pressure ratio balance between the fan and core compression systems.

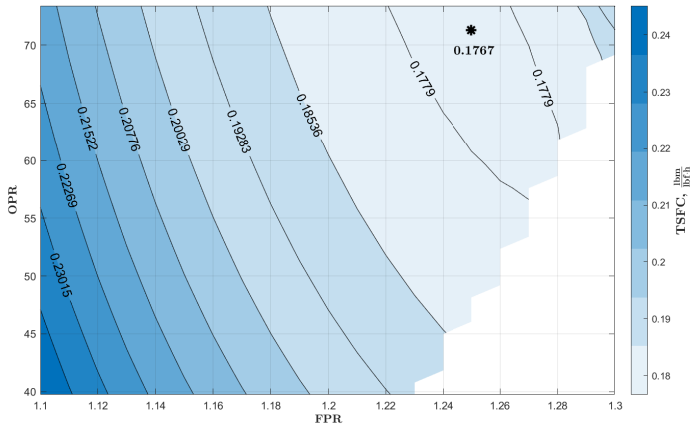

Fig. 2.11 Hydrogen-fuelled engine with preheated fuel: TSFC as a function of OPR and FPR with fixed BPR = 16.2 and TIT = 1600 K.

The minimum TSFC for this configuration is 0.1767 lbm/(lbf·h), approximately 35% of the kerosene baseline value. The required hydrogen mass flow rate is 0.199 kg/s. Comparing TSEC values:

$$\frac{\mathrm{TSEC}_{\mathrm{H_2}}}{\mathrm{TSEC}_{\mathrm{kerosene}}} = \frac{(\dot{m}_{\mathrm{fuel}} \cdot \mathrm{LHV})_{\mathrm{H_2}}}{(\dot{m}_{\mathrm{fuel}} \cdot \mathrm{LHV})_{\mathrm{kerosene}}} \approx 0.98 \tag{2.15}$$

The hydrogen configuration achieves approximately 2% lower energy consumption per unit thrust compared to kerosene at cruise conditions, consistent with findings in Ref. [2]. This modest improvement reflects the competing effects of hydrogen's high gravimetric energy density, favourable turbine exhaust properties, and the energy penalty of heating cryogenic fuel (in this idealized case, assumed to occur without affecting engine performance).

2.5.4 Regenerative Fuel Heating Cycle

The regenerative fuel heating cycle exploits exhaust gas thermal energy to warm the hydrogen fuel before combustion, recovering waste heat that would otherwise be rejected through the nozzle. Figure 2.12 illustrates the cycle architecture. Liquid hydrogen at approximately 50 K exits the cryogenic pump at station 21 and flows through a heat exchanger positioned downstream of the low-pressure turbine (station 5). The heated fuel at station 22 then enters the combustor. The cooled exhaust gases at station 5.5 proceed to the nozzle.

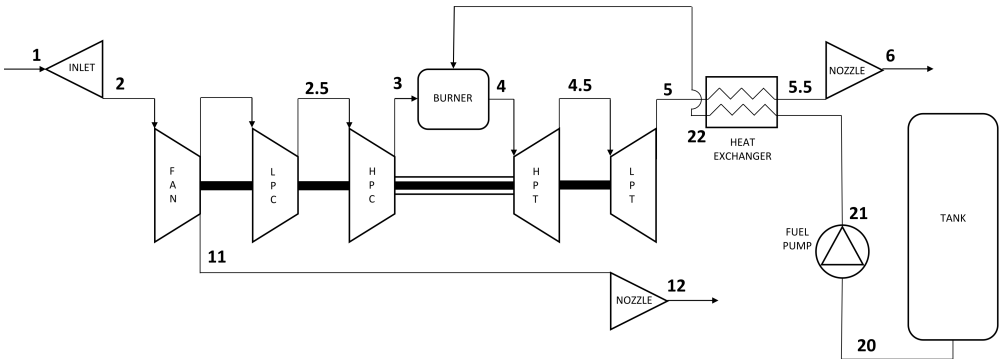

Fig. 2.12 Schematic diagram of the regenerative fuel heating cycle. Hydrogen is heated by exhaust gases downstream of the low-pressure turbine before entering the combustor.

The heat exchanger is modelled using an effectiveness-NTU approach with effectiveness $\varepsilon_{\mathrm{HX}} = 0.8$ [23]. The outlet temperatures are determined from:

$$C = \dot{m} c_p, \qquad C_{\min} = \min(C_{\mathrm{H_2}}, C_{\mathrm{air}}) \tag{2.16}$$

$$T_{t,22} = T_{t,21} + \varepsilon_{\mathrm{HX}} (T_{t,5} - T_{t,21}) \frac{C_{\min}}{C_{\mathrm{H_2}}} \tag{2.17}$$

$$T_{t,5.5} = T_{t,5} - \varepsilon_{\mathrm{HX}}(T_{t,5} - T_{t,21})\frac{C_{\min}}{C_{\mathrm{air}}} \tag{2.18}$$

where $T_{t,5}$ and $T_{t,5.5}$ are the total temperatures at the air inlet and outlet of the heat exchanger, and $T_{t,21}$ and $T_{t,22}$ are the corresponding fuel temperatures. The specific heat values are evaluated at total conditions; for hydrogen in the supercritical state (pressure above critical pressure of 1.3 MPa at temperatures below the critical temperature of 33 K), phase change does not occur within the heat exchanger, simplifying the thermal analysis. A 4% total pressure drop is applied to the air stream across the heat exchanger [23].

The fuel pump configuration follows Ref. [23]: a two-stage centrifugal pump mechanically connected to the low-pressure spool via the engine gearbox. The rotational speed is fixed by the gearbox ratio and low-pressure spool speed. Pump efficiency is estimated from specific speed correlations in Ref. [107], with the specific speed defined in Eq. (2.12).

Repeating the parametric optimization for this cycle yields the TSFC surface shown in Figure 2.13. The minimum TSFC is 0.1781 lbm/(lbf·h), occurring at FPR = 1.25 and OPR = 66.4, with hydrogen mass flow rate of 0.200 kg/s.

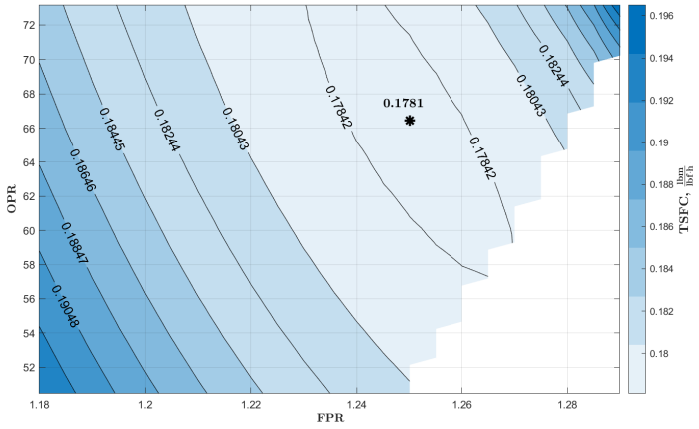

Fig. 2.13 Regenerative fuel heating cycle: TSFC as a function of OPR and FPR with fixed BPR = 16.2 and TIT = 1600 K.

Comparing this result to the idealized preheated hydrogen case (Section 2.5.3) reveals a slight TSFC increase of approximately 0.8%. This apparent contradiction with Brewer's assertion [23] that fuel preheating improves TSFC requires clarification. In the idealized case, hydrogen was assumed to enter the combustor at

298 K via an external and unmodelled heat source, representing an unrealistic upper performance bound. In contrast, the regenerative fuel heating cycle accounts for the actual thermal energy required to vaporize and heat the cryogenic fuel, extracting this energy from the engine's own exhaust stream. The heat exchanger also introduces a pressure loss in the main gas path, further penalizing performance. Nevertheless, the regenerative cycle TSFC remains approximately 36% of the kerosene baseline, and the TSEC is approximately 1% lower than kerosene, demonstrating that effective thermal integration of cryogenic hydrogen can achieve performance comparable to conventional fuels despite the challenges of fuel conditioning.

2.5.5 Expander Cycle

The expander cycle extends the regenerative fuel heating concept by routing the heated hydrogen through a dedicated expansion turbine before injection into the combustor, as shown in Figure 2.14. This turbine (station 23) generates auxiliary power for aircraft systems (assumed to be 125 hp following Ref. [23]), eliminating or reducing the extraction of mechanical power from the engine's main shaft through the accessory gearbox.

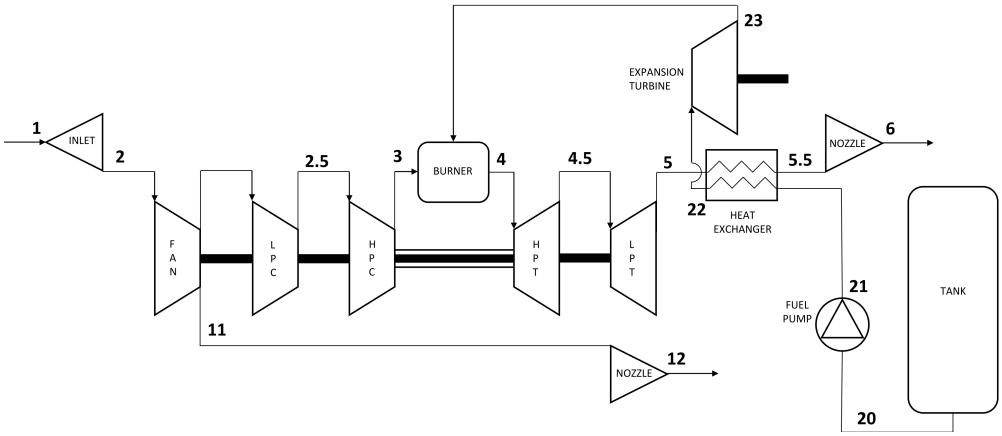

Fig. 2.14 Schematic diagram of the expander cycle. Heated hydrogen expands through an auxiliary turbine (station 23) before entering the combustor, generating power for aircraft systems.

To maintain adequate fuel injection pressure at the combustor inlet after expansion, the pump must pressurize the hydrogen to a higher level than in the regenerative fuel heating cycle. This requirement has two consequences: first, the pump extracts

additional work from the low-pressure spool to which it is connected. Second, the higher discharge pressure reduces the pump's specific speed (Eq. 2.12), resulting in lower hydraulic efficiency according to the correlations in Ref. [107]. The hydrogen also enters the heat exchanger at elevated pressure, altering its thermodynamic properties including specific heat. Table 2.2 presents the total pressure and temperature at key stations for both the regenerative fuel heating and expander cycles.

Table 2.2 Total pressure and temperature at engine stations for the regenerative fuel heating and expander cycles.

	Regenerative Fuel Heating		Expander Cycle	
Station	p_t (kPa)	T_t (K)	p_t (kPa)	T_t (K)
1	56.5	265.8	56.5	265.8
2	56.3	265.8	56.3	265.8
11	70.3	285.4	70.3	285.4
3	3755.6	962.2	3755.6	962.2
4	3586.7	1600.0	3586.7	1600.0
5	61.7	650.7	61.6	650.7
5.5	59.2	599.3	59.1	597.8
21	3912.1	50.0	4864.3	50.0
22	3755.6	530.9	4669.8	530.9
23	—	—	3755.6	499.1

Despite these differences, the integrated effects on engine performance are minimal. At fixed FPR, OPR, BPR, TIT, and thrust requirement, the fuel mass flow rate changes by less than 1% between the regenerative fuel heating and expander cycles. Consequently, the TSFC values are nearly identical, both approximately 36% of the kerosene baseline. This result confirms Brewer's observation [23] that the two cycles exhibit similar propulsion performance. The expander cycle offers the advantage of decoupled auxiliary power generation without extracting shaft power from the main engine, but introduces greater mechanical complexity with the addition of a fuel turbine and associated control systems.

2.5.6 Comparative Performance Assessment

Table 2.3 summarizes the key performance metrics for all configurations analysed. The TSFC values demonstrate that hydrogen-fuelled engines, even with realistic

thermal integration of cryogenic fuel, achieve fuel consumption approximately one-third that of kerosene engines when compared on a mass basis. When accounting for the different heating values through the TSEC metric, hydrogen configurations achieve 1–2% lower energy consumption per unit thrust at cruise conditions under the assumptions and simplifications of this analysis.

Table 2.3 Performance comparison of propulsion system configurations at cruise.

Configuration	FPR	OPR	TSFC (lbm/(lbf·h))	$\dot{m}_{\text{fuel}}$ (kg/s)	TSEC Ratio (vs. kerosene)
Kerosene baseline	1.23	67.5	0.4995	0.562	1.00
H_2 preheated (ideal)	1.25	71.3	0.1767	0.199	0.98
Regenerative fuel heating	1.25	66.4	0.1781	0.200	0.99
Expander cycle	1.25	66.2	0.1784	0.201	0.99

Several observations emerge from this comparative assessment:

1. **Fuel mass reduction**: Hydrogen's high gravimetric energy density enables a 64% reduction in fuel mass flow rate compared to kerosene, with corresponding implications for aircraft weight and combustion chamber size.

2. **Optimal pressure ratios**: The hydrogen configurations favour slightly lower OPR than kerosene (66.2-66.4 versus 67.5), leaving aside the ideal case, reflecting the altered thermodynamic properties of water-rich combustion products. The optimal FPR increases slightly from 1.23 to 1.25.

3. **Thermal integration penalty**: The regenerative fuel heating and expander cycles incur a modest TSFC penalty (approximately 0.8–1%) compared to the idealized preheated hydrogen case, primarily due to heat exchanger pressure losses in the exhaust stream. However, there is an overall advantage over kerosene.

4. **Cycle equivalence**: The regenerative fuel heating and expander cycles exhibit nearly identical TSFC performance, confirming that the choice between them should be driven by considerations other than cruise efficiency, such as auxiliary power requirements, system complexity, and off-design operability.

These results should be interpreted within the context of the analysis assumptions: steady-state cruise operation, fixed component efficiencies independent of fuel type,

perfect mixing and combustion, and neglect of installation effects. The modest TSEC advantage of hydrogen (1–2%) is consistent with the range of outcomes reported in the literature [2], which spans from net energy penalties to moderate benefits depending on aircraft configuration, mission profile, and technology assumptions. The present analysis demonstrates that with appropriate thermal integration, hydrogen propulsion can achieve energy consumption comparable to or slightly better than kerosene at cruise, while offering the environmental benefits of zero direct CO_2 emissions and the potential for substantial NO_x reduction through ultra-lean combustion.

2.6 Fuel System Integration and Storage Feasibility

The preceding section established the propulsion system performance characteristics of hydrogen turbofan engines. However, the practical implementation of hydrogen aviation depends critically on the fuel storage system. This section examines the feasibility of integrating liquid hydrogen tanks within the reference BWB aircraft, assessing tank sizing, gravimetric efficiency, and boil-off behaviour.

2.6.1 Liquid Hydrogen Tank Design

While hydrogen possesses higher energy density by weight compared to kerosene (120 MJ/kg versus 43 MJ/kg), its substantially lower density results in approximately four times the volume for equivalent energy storage [2]. This volumetric penalty necessitates careful tank design and placement within the airframe.

Two principal tank architectures are employed in aircraft applications: integral tanks, which form part of the airframe structure and share structural loads, and non-integral tanks, which are independent pressure vessels designed to withstand only internal forces. Integral tanks offer superior volumetric efficiency by exploiting available airframe volume, but require detailed knowledge of structural design and load paths. For the present feasibility assessment, non-integral tanks are selected, allowing generalized analysis without requiring complete airframe structural integration. This approach is conservative from a weight perspective but suitable for demonstrating storage feasibility.

The BWB configuration of Ref. [71] provides substantial internal volume within the thick wing centre section, accommodating large cylindrical tanks. The tank geometry selected comprises horizontal cylinders with spherical end caps, a configuration that balances structural efficiency with manufacturability and integration constraints.

2.6.2 Tank Sizing and Configuration

The total hydrogen mass required for the mission is determined from the engine fuel consumption rates calculated in Section 2.5. For the three-engine configuration analysed (regenerative fuel heating cycle), each engine consumes approximately 0.200 kg/s at cruise. Over a 12-hour mission with varying power settings, the total fuel requirement is estimated at 26,550 kg of liquid hydrogen.

Several factors influence the required tank volume beyond the liquid hydrogen mass alone. First, a portion of the tank volume must be allocated for gaseous hydrogen to facilitate pressure management and venting. Following Ref. [158], 3% of the total tank volume is reserved for gaseous hydrogen (GH_2), with the remaining 97% available for liquid storage. Second, the storage density depends on the operating pressure. Higher venting pressures reduce liquid density while increasing gas density, resulting in lower overall storage density and larger required tank volumes. The minimum operating pressure is set at 0.8 bar, above maximum ambient pressure to prevent air ingress. The maximum operating pressure (venting pressure) is set at 1.44 bar, following the recommendation in Ref. [158] to minimize venting pressure and thereby reduce both tank structural mass and required volume. Storage temperature is maintained at 20 K through insulation, with initial pressure of 1.2 bar slightly above the saturation pressure at this temperature.

At these conditions, liquid hydrogen density is 71.28 kg/m^3, yielding a required liquid volume of 373 m^3 for 26,550 kg of fuel. To accommodate the tapered geometry of the BWB wing, six tanks of three different sizes are employed, consistent with the configuration in Ref. [71]. Larger tanks are positioned near the thick inner wing section, with progressively smaller tanks towards the thinner outer wing.

The tank dimensions are constrained by the fuselage height $H_{\text{ref}} = 4$ m, which limits the maximum tank diameter for the innermost tanks adjacent to the cabin. Tank lengths are expressed relative to the cabin root chord $C_0 = 43.5$ m. Table 2.4

presents the normalized dimensions and volume contributions for the three tank sizes (dimensions shown for one semi-wing; the configuration is symmetric about the aircraft centreline).

Table 2.4 Normalized tank dimensions and volume contributions for one semi-wing.

Tank	D_i / H_{ref}	D_o / H_{ref}	L_{tank} / C_0	V_{tank} / V_{LH_2}
Tank 1	0.625	0.675	0.437	20.5%
Tank 2	0.575	0.625	0.368	16.9%
Tank 3	0.525	0.575	0.328	12.6%

2.6.3 Tank Structure and Gravimetric Efficiency

The tank structural design follows a composite sandwich configuration based on Ref. [9]: an inner aluminium liner (Al 2219) for hydrogen compatibility, a middle foam insulation layer, and an outer carbon fibre reinforced plastic (CFRP) shell for structural strength and environmental protection. The aluminium liner prevents hydrogen embrittlement of structural materials, while the CFRP layer minimizes thermal expansion coefficient mismatches and provides the primary load-bearing capability.

Layer thicknesses are specified as: aluminium liner $t_{\text{Al}} = 0.01$ m, foam insulation $t_{\text{foam}} = 0.18$ m, CFRP shell $t_{\text{CFRP}} = 0.01$ m, yielding a total thickness-to-diameter ratio of approximately 0.08. Using material densities for each layer and the tank geometries from Table 2.4, the total tank system mass (for all six tanks) is calculated as 28,115 kg.

The gravimetric efficiency η_{grav} quantifies the fraction of the total fuel system mass attributable to usable fuel:

$$\eta_{\text{grav}} = \frac{m_{\text{LH}_2}}{m_{\text{LH}_2} + m_{\text{tank}}} = \frac{26,550}{26,550 + 28,115} = 0.485 \tag{2.19}$$

This value of 48.5% compares favourably with the 48.8% gravimetric efficiency reported in Ref. [71] for a similar BWB configuration, confirming that the non-integral tank approach adopted here achieves mass characteristics consistent with more integrated designs. Gravimetric efficiencies above 40% are generally consid-

ered necessary for practical hydrogen aircraft [23]; the present design exceeds this threshold with margin.

2.6.4 Boil-Off Assessment

Even with substantial insulation, heat leakage from the external environment causes gradual evaporation of liquid hydrogen, known as boil-off. If the evaporated gas is not vented, tank pressure increases, potentially exceeding safe operating limits. The 3% gas volume allocation (Section 2.6.2) provides capacity for pressure management, but must be verified as sufficient across the mission duration.

Boil-off behaviour is simulated using a two-phase lumped thermodynamic model accurately described in [164]. The model solves coupled ordinary differential equations for three control volumes: ullage (gas space), bulk liquid, and tank walls. Mass and energy balances govern the thermodynamic evolution of each control volume, accounting for heat transfer at interfaces, evaporation at the liquid-gas interface, and mass withdrawal as fuel is supplied to the engines. Heat transfer coefficients are determined from empirical correlations, with hydrogen properties obtained from the CoolProp database [16].

Heat ingress from the external environment is estimated following Ref. [9]:

$$\dot{Q} = \frac{\Delta T}{R_{\text{th}}} A, \qquad R_{\text{th}} = \sum_i \frac{t_i}{k_i} \tag{2.20}$$

where ΔT is the temperature difference between ambient conditions and the tank interior, A is the tank surface area, R_{th} is the thermal resistance of the tank walls, and the summation is over the three layers (aluminium, foam, CFRP) with thicknesses t_i and thermal conductivities k_i.

The simulation is initialized with the tank at 20 K and 1.2 bar. As fuel is withdrawn during the mission (proportionally from each tank according to its capacity to prevent premature depletion of smaller tanks), heat ingress causes ullage pressure to rise. When pressure reaches the venting limit (1.44 bar), excess gas is released. The simulation proceeds over the 12-hour mission duration to determine the accumulated boil-off mass.

Results indicate that the boil-off mass remains within the 3% gas allocation for all three tank sizes: approximately 1.5% for Tank 1 (largest), 1.6% for Tank 2, and 1.8% for Tank 3 (smallest). The trend of higher boil-off percentage in smaller tanks reflects their larger surface-area-to-volume ratio, which increases the relative heat ingress per unit stored hydrogen. These values confirm that the selected tank configuration and insulation thickness provide adequate boil-off management over the mission duration without requiring active cooling or excessive venting losses.

2.6.5 Integration Feasibility Summary

The fuel system analysis demonstrates that liquid hydrogen storage is feasible for the reference BWB aircraft configuration under the mission requirements examined. The six-tank arrangement accommodates the required fuel mass within the geometric constraints of the thick wing centre section, achieving a gravimetric efficiency of 48.5% comparable to more integrated designs. Boil-off remains manageable with passive insulation, with losses below 2% of stored fuel over the 12-hour mission.

Several caveats accompany these conclusions. First, the analysis assumes non-integral tanks, which are conservative from a weight perspective; integral tanks exploiting airframe structure could achieve higher gravimetric efficiency. Second, the boil-off assessment considers cruise and assumes uniform heat ingress rates; ground hold periods before takeoff and descent/landing phases may experience different thermal environments requiring additional analysis. Third, the present study considers only steady-state cruise conditions for the propulsion system; transient operations during takeoff, climb, and descent involve different fuel flow rates and thermal conditions that must be examined in future work. Nevertheless, the results establish that hydrogen fuel system integration can meet the performance targets of long-range BWB aircraft while maintaining reasonable system mass characteristics.

2.7 Summary and Transition to Detailed Analysis

This chapter has provided both the fundamental physics and system-level context for hydrogen combustion in aviation propulsion. The key findings can be summarized as follows:

- **Fundamental properties**: The laminar flame speed of hydrogen exceeds that of hydrocarbon fuels by nearly an order of magnitude across the range of conditions relevant to gas turbine combustors. The sub-unity Lewis number ($Le \approx 0.3$) leads to thermo-diffusive instabilities that fundamentally alter flame structure compared to unity-Lewis-number hydrocarbon flames.

- **NO_x formation pathways**: Nitrogen oxide formation in hydrogen flames proceeds through thermal (Zeldovich), N_2O intermediate, and NNH pathways, with relative contributions depending on local temperature, pressure, and mixture composition. Water/steam injection offers substantial NO_x reduction potential through combined thermal, dilution, and chemical effects.

- **System-level opportunities**: Integrated propulsion system modelling quantifies the performance implications of alternative thermodynamic cycles exploiting hydrogen's cryogenic properties. Regenerative fuel heating and expander cycles achieve TSFC approximately 36% of kerosene baseline values, corresponding to TSEC 1% lower than kerosene at cruise conditions.

- **Storage feasibility**: Liquid hydrogen fuel systems can be integrated within BWB airframe configurations with gravimetric efficiencies exceeding 48%, with passive insulation adequate for boil-off management over 12-hour missions.

The Cantera-based property calculations and NPSS system modelling presented in this chapter establish the opportunity space for hydrogen aviation: the potential benefits are substantial, spanning environmental, operational, and performance domains. However, realization of these benefits depends critically on combustor technologies that can manage hydrogen's challenging combustion characteristics, particularly flashback risk and NO_x formation, while maintaining the reliability and durability required for commercial aviation.

Chapter 3 addresses this challenge by developing and validating computational methods for the analysis of hydrogen combustor designs. The micromix injector study examines a flashback prevention strategy, while the multi-fidelity modelling framework enables efficient design space exploration with limited experimental data.

Chapter 3

CFD Simulations of Hydrogen Flames

The design of hydrogen-fuelled gas turbine combustors requires an understanding of the interaction between turbulent mixing, chemical kinetics, and heat release. While experimental campaigns provide essential validation data, they are often limited by optical accessibility and the safety constraints associated with high-pressure hydrogen testing. Computational Fluid Dynamics (CFD) serves as a pivotal tool in this context, enabling the detailed resolution of flow field structures to navigate the preliminary design landscape.

This chapter presents the numerical investigation of hydrogen flames applied to two distinct combustion concepts, employing a range of modelling fidelities suited to the specific physical phenomena under scrutiny.

The chapter opens with the numerical framework in Sec. 3.1, which unifies the governing equations for the simulations presented thereafter. This section details the Reynolds-Averaged Navier-Stokes (RANS) formulations in both three-dimensional and two-dimensional axisymmetric coordinate systems, the Large Eddy Simulation (LES) approach, used for validation, and the models used for turbulence-chemistry interactions.

In Sec. 3.2, the investigation focuses on hydrogen micromix combustors, a technology designed to minimize NO_x emissions through the generation of miniaturized diffusion flamelets. Here, the primary objective is the evaluation of pollutant formation mechanisms. Consequently, the study utilizes three-dimensional RANS simulations coupled with a decoupled post-processing approach to quantify NO and

N_2O emissions. This section specifically assesses the impact of converging nozzle geometries on flame structure and emissions production [38].

Subsequently, Sec. 3.3 addresses the AHEAD lean premixed swirl-stabilized burner, with a primary focus on aerodynamic stability and flashback prevention via flame position tracking. Recognizing that mapping stability limits requires extensive parametric exploration, this section moves beyond computationally expensive 3D simulations. Instead, it proposes a multi-fidelity modelling strategy that leverages cost-effective 2D axisymmetric RANS simulations corrected by sparse experimental data [41]. The validity of the 2D approximation is critically assessed against experimental PIV data and high-fidelity 3D LES to ensure the relevant flow physics are captured.

3.1 Numerical Framework

This section presents the numerical framework employed for the CFD simulations described in the subsequent sections. The governing equations and turbulence-chemistry interaction models are introduced in a unified manner to avoid repetition in later discussions.

3.1.1 RANS: Three-Dimensional Formulation

The simulations employ the RANS equations for steady-state, three-dimensional, compressible, reacting flow. Introducing the Favre-filtered variables, denoted by a tilde, obtained for a generic scalar β as:

$$\tilde{\beta} = \frac{\overline{\rho\beta}}{\bar{\rho}} \tag{3.1}$$

where the overbar indicates the typical Reynolds averaging and ρ represents the density. The corresponding Navier-Stokes equations are written in compressible form as follows [6]:

$$\frac{\partial \overline{\rho}}{\partial t} + \nabla \cdot \left(\overline{\rho}\vec{\tilde{v}}\right) = 0 \tag{3.2}$$

$$\frac{\partial\left(\overline{\rho}\,\vec{\tilde{v}}\right)}{\partial t}+\nabla\cdot\left(\overline{\rho}\vec{\tilde{v}}\vec{\tilde{v}}\right)=-\nabla\overline{p}+\nabla\cdot\left(\hat{\tau}\right) \tag{3.3}$$

$$\frac{\partial\left(\overline{\rho}\left(\tilde{e}+\frac{\tilde{v}^2}{2}\right)\right)}{\partial t}+\nabla\cdot\left(\overline{\rho}\vec{\tilde{v}}\left(\tilde{h}+\frac{\tilde{v}^2}{2}\right)\right)= \nabla\cdot\left(k_{eff}\nabla\tilde{T}-\sum_j\tilde{h}_j\vec{J}_j+\hat{\tau}\cdot\vec{\tilde{v}}\right)+\Omega \tag{3.4}$$

In Eqs. (3.2)–(3.4), p represents the pressure, $\hat{\tau}$ the total stress tensor, e the specific internal energy, h the specific sensible enthalpy, k_{eff} the effective conductivity, J_j the diffusion flux of species j, and Ω the volumetric heat of chemical reactions. The total stress tensor $\hat{\tau}$ is formed by the sum of the laminar and turbulent stresses:

$$\hat{\tau}_{ij}=2\left(\mu+\mu_t\right)\left(\tilde{S}_{ij}-\frac{1}{3}\frac{\partial\tilde{v}_k}{\partial x_k}\delta_{ij}\right)-\frac{2}{3}\bar{\rho}\,\kappa\delta_{ij} \tag{3.5}$$

with μ and μ_t indicating the laminar and turbulent dynamic viscosity, respectively, κ the turbulent kinetic energy and S_{ij} the strain rate tensor:

$$\tilde{S}_{ij}=\frac{1}{2}\left(\frac{\partial\tilde{v}_i}{\partial x_j}+\frac{\partial\tilde{v}_j}{\partial x_i}\right) \tag{3.6}$$

Turbulence closure is provided by the κ-ω Shear Stress Transport (SST) model [102], which combines the advantages of the κ-ω formulation in the near-wall region with the κ-ε behaviour in the freestream, providing improved predictions for flows with adverse pressure gradients and separation.

The effective thermal conductivity k_{eff} is defined as:

$$k_{eff}=k+\frac{c_p\mu_t}{Pr_t} \tag{3.7}$$

where k is the laminar thermal conductivity, c_p is the specific heat at constant pressure and Pr_t is the turbulent Prandtl number. In this work, Pr_t was selected to be

0.85, following Ref. [72]. The c_p is obtained as the mixture-average of the c_p of all the species in the mixture, while the molar heat capacity at constant pressure $\hat{c}_{p,j}$ is approximated by NASA7 polynomials [100] for the single species j:

$$\frac{\hat{c}_{p,j}(\tilde{T})}{R} = a_{0,j} + a_{1,j}\tilde{T} + a_{2,j}\tilde{T}^2 + a_{3,j}\tilde{T}^3 + a_{4,j}\tilde{T}^4 \tag{3.8}$$

with R being the universal gas constant. Thus, the sensible enthalpy for the single species j in the case of ideal gas approximation is calculated as:

$$\tilde{h}_j = \int_{T_{ref}}^{\tilde{T}} c_{p,j}\, d\tilde{T} \tag{3.9}$$

with T_{ref} being 298.15 K. Then, the sensible enthalpy of the mixture is evaluated as:

$$\tilde{h} = \sum_j \tilde{Y}_j \tilde{h}_j \tag{3.10}$$

The diffusion flux of species j is defined, including the full multicomponent and Soret diffusion [6], as:

$$\vec{J}_j = -\rho \sum_{k=1}^{M-1} D_{j,k} \nabla \tilde{Y}_k - \frac{\mu_t}{Sc_t} \nabla \tilde{Y}_j - D_{\tilde{T},j} \frac{\nabla \tilde{T}}{\tilde{T}} \tag{3.11}$$

where Sc_t is the turbulent Schmidt number, Y is the mass fraction, $D_{j,k}$ are the generalized Fick's law diffusion coefficients obtained from the Maxwell–Stefan equations [6] and $D_{T,j}$ is the Soret diffusion coefficient.

Finally, the volumetric heat of reaction Ω is calculated as:

$$\Omega = -\sum_j \frac{h_j^0}{\mathscr{M}_j} \mathscr{R}_j \tag{3.12}$$

where h_j^0, $\mathscr{M}_j$ and $\mathscr{R}_j$ are the enthalpy of formation, the molecular weight and the volumetric rate of creation, respectively, of the j^{th} species.

In addition to Eqs. (3.2)–(3.4), a conservation equation is considered for $M-1$ species, where M denotes the total number of species in the mixture. The mass fraction of the excluded species is calculated as 1 minus the sum of the mass fractions of all the other species. To minimize the numerical error, N_2 is chosen to be the

species not directly computed by a transport equation, as it is expected to have the largest mass fraction. These conservation equations are then written as:

$$\frac{\partial\left(\overline{\rho}\,\tilde{Y}_j\right)}{\partial t}+\nabla\cdot\left(\overline{\rho}\vec{\tilde{v}}\tilde{Y}_j\right)=-\nabla\cdot\vec{J}_j+\mathscr{R}_j \tag{3.13}$$

3.1.2 RANS: Two-Dimensional Axisymmetric with Swirl

For certain configurations, computational efficiency is achieved by employing a 2D steady swirl axisymmetric flow approximation. This formulation neglects every derivative along the azimuthal θ direction while accounting for a tangential component of velocity v_θ variable along the axial x and radial r directions [1, 6]:

$$\frac{\partial(\cdot)}{\partial\theta}=0 \tag{3.14}$$

$$\vec{v}(x,r)=[v_x,v_r,v_\theta] \tag{3.15}$$

The analysis takes into account variable density using Favre-averaged variables, as defined previously. The resulting compressible 2D RANS equations are written in cylindrical coordinates as:

$$\frac{\partial\bar{\rho}}{\partial t}+\frac{1}{r}\frac{\partial}{\partial x}\left(\bar{\rho}r\tilde{v}_x\right)+\frac{\partial}{\partial r}\left(\bar{\rho}\tilde{v}_r\right)+\bar{\rho}\frac{\tilde{v}_r}{r}=0 \tag{3.16}$$

$$\begin{aligned}\frac{\partial}{\partial t}\left(\bar{\rho}\tilde{v}_x\right)+\frac{1}{r}\frac{\partial}{\partial x}\left(r\bar{\rho}\tilde{v}_x\tilde{v}_x\right)+\frac{1}{r}\frac{\partial}{\partial r}\left(r\bar{\rho}\tilde{v}_r\tilde{v}_x\right)=&\\ -\frac{\partial\bar{p}}{\partial x}+\frac{1}{r}\frac{\partial}{\partial x}\left(r\hat{\tau}_{xx}\right)+\frac{1}{r}\frac{\partial}{\partial r}\left(r\hat{\tau}_{rx}\right)&\end{aligned} \tag{3.17}$$

$$\frac{\partial}{\partial t}(\bar{\rho}\tilde{v}_r)+\frac{1}{r}\frac{\partial}{\partial x}(r\bar{\rho}\tilde{v}_x\tilde{v}_r)+\frac{1}{r}\frac{\partial}{\partial r}(r\bar{\rho}\tilde{v}_r\tilde{v}_r)= -\frac{\partial \bar{p}}{\partial r}+\frac{1}{r}\frac{\partial}{\partial x}(r\hat{\tau}_{xr})+\frac{1}{r}\frac{\partial}{\partial r}(r\hat{\tau}_{rr})-\frac{1}{r}\hat{\tau}_{\theta\theta}+\bar{\rho}\frac{\tilde{v}_\theta^2}{r} \quad (3.18)$$

$$\frac{\partial}{\partial t}(\bar{\rho}\tilde{v}_\theta)+\frac{1}{r}\frac{\partial}{\partial x}(r\bar{\rho}\tilde{v}_x\tilde{v}_\theta)+\frac{1}{r}\frac{\partial}{\partial r}(r\bar{\rho}\tilde{v}_r\tilde{v}_\theta)= \frac{1}{r}\frac{\partial}{\partial x}(r\hat{\tau}_{x\theta})+\frac{1}{r}\frac{\partial}{\partial r}(r\hat{\tau}_{r\theta})+\frac{1}{r}\hat{\tau}_{r\theta}-\bar{\rho}\frac{\tilde{v}_r\tilde{v}_\theta}{r} \quad (3.19)$$

In these equations $\vec{\mathbf{v}} = [v_x, v_r, v_\theta]$ represents the velocity vector with axial, radial, and tangential components. The variable r represents the radius, p represents the pressure and $\hat{\tau}$ represents the total stress tensor. The total stress tensor is a sum of the viscous and turbulent stress tensors, and its components are:

$$\hat{\tau}_{xx} = (\mu+\mu_t)\left(2\frac{\partial \tilde{v}_x}{\partial x}-\frac{2}{3}\nabla\cdot\vec{\tilde{v}}\right)-\frac{2}{3}\bar{\rho}\kappa \quad (3.20)$$

$$\hat{\tau}_{rr} = (\mu+\mu_t)\left(2\frac{\partial \tilde{v}_r}{\partial r}-\frac{2}{3}\nabla\cdot\vec{\tilde{v}}\right)-\frac{2}{3}\bar{\rho}\kappa \quad (3.21)$$

$$\hat{\tau}_{\theta\theta} = (\mu+\mu_t)\left(2\frac{\tilde{v}_r}{r}-\frac{2}{3}\nabla\cdot\vec{\tilde{v}}\right)-\frac{2}{3}\bar{\rho}\kappa \quad (3.22)$$

$$\hat{\tau}_{xr} = (\mu+\mu_t)\left(\frac{\partial \tilde{v}_x}{\partial r}+\frac{\partial \tilde{v}_r}{\partial x}\right) \quad (3.23)$$

$$\hat{\tau}_{x\theta} = (\mu+\mu_t)\left(\frac{\partial \tilde{v}_\theta}{\partial x}\right) \quad (3.24)$$

$$\hat{\tau}_{r\theta} = (\mu+\mu_t)\left(\frac{\partial \tilde{v}_\theta}{\partial r}-\frac{\tilde{v}_\theta}{r}\right) \quad (3.25)$$

with the divergence of the velocity field that in 2D axisymmetric condition is expressed as:

$$\nabla\cdot\vec{\tilde{v}} = \frac{\partial \tilde{v}_x}{\partial x}+\frac{\partial \tilde{v}_r}{\partial r}+\frac{\tilde{v}_r}{r} \quad (3.26)$$

As the transport of species in the flow domain is considered, the energy equation contains, on top of the terms accounting for energy transfer due to conduction A

and viscous dissipation B, a term accounting for the diffusion of the enthalpy of the species C and a term accounting for the heat generation from chemical reactions Ω. In 2D axisymmetric RANS, the energy equation is thus written as:

$$\frac{\partial}{\partial t}\left(\bar{\rho}\tilde{E}\right)+\frac{1}{r}\frac{\partial}{\partial x}\left(\bar{\rho}r\tilde{H}\tilde{v}_x\right)+\frac{1}{r}\frac{\partial}{\partial r}\left(\bar{\rho}r\tilde{H}\tilde{v}_r\right)=\frac{1}{r}\left(+A+B-C\right)+\Omega \tag{3.27}$$

with:

$$A=\frac{\partial}{\partial r}\left(rk_{eff}\frac{\partial\tilde{T}}{\partial r}\right)+\frac{\partial}{\partial x}\left(rk_{eff}\frac{\partial\tilde{T}}{\partial x}\right) \tag{3.28}$$

$$B=\frac{\partial}{\partial x}\left(r\left(\hat{\tau}_{xx}\tilde{v}_x+\hat{\tau}_{xr}\tilde{v}_r+\hat{\tau}_{x\theta}\tilde{v}_\theta\right)\right)+\frac{\partial}{\partial r}\left(r\left(\hat{\tau}_{xr}\tilde{v}_x+\hat{\tau}_{rr}\tilde{v}_r+\hat{\tau}_{r\theta}\tilde{v}_\theta\right)\right) \tag{3.29}$$

$$C=\frac{\partial}{\partial x}\left(r\sum_j\tilde{h}_j\vec{J}_j\cdot\hat{x}\right)+\frac{\partial}{\partial r}\left(r\sum_j\tilde{h}_j\vec{J}_j\cdot\hat{r}\right) \tag{3.30}$$

$$\Omega=-\sum_j\frac{h_j^0}{\mathscr{M}_j}\mathscr{R}_j \tag{3.31}$$

where E is the sum of internal and kinetic energy, $\vec{J}_j$ is the diffusion flux of species j, and h_j^0, $\mathscr{M}_j$ and $\mathscr{R}_j$ are the enthalpy of formation, the molecular weight and the volumetric rate of creation, respectively, of the j^{th} species. The variable $\tilde{H}$ represents:

$$\tilde{H}=\tilde{h}+\frac{\tilde{v}^2}{2} \tag{3.32}$$

The diffusion flux of species j in the 2D formulation is defined, including the full multicomponent and Soret diffusion [6], as:

$$\begin{aligned}\vec{J}_j=-\rho\sum_{k=1}^{M-1}D_{j,k}\left(\frac{\partial\tilde{Y}_k}{\partial r}\hat{r}+\frac{\partial\tilde{Y}_k}{\partial x}\hat{x}\right)+\\ -\frac{\mu_t}{Sc_t}\left(\frac{\partial\tilde{Y}_j}{\partial r}\hat{r}+\frac{\partial\tilde{Y}_j}{\partial x}\hat{x}\right)-\frac{D_{\tilde{T},j}}{\tilde{T}}\left(\frac{\partial\tilde{T}}{\partial r}\hat{r}+\frac{\partial\tilde{T}}{\partial x}\hat{x}\right)\end{aligned} \tag{3.33}$$

where Sc_t is the turbulent Schmidt number, equal to 0.7. The unit vectors $\hat{x}$ and $\hat{r}$ denote the axial and radial directions, respectively.

A conservation equation is considered for $M-1$ species, where M denotes the total number of species in the mixture, with the mass fraction of N_2 determined by subtraction. These conservation equations are written as:

$$\frac{\partial\left(\bar{\rho}\tilde{Y}_j\right)}{\partial t}+\frac{1}{r}\frac{\partial\left(r\bar{\rho}\tilde{v}_x\tilde{Y}_j\right)}{\partial x}+\frac{1}{r}\frac{\partial\left(r\bar{\rho}\tilde{v}_r\tilde{Y}_j\right)}{\partial r}= -\frac{1}{r}\left[\frac{\partial}{\partial x}\left(r\vec{J}_j\cdot\hat{x}\right)+\frac{\partial}{\partial r}\left(r\vec{J}_j\cdot\hat{r}\right)\right]+\mathscr{R}_j \quad (3.34)$$

3.1.3 Turbulence-Chemistry Interaction: Eddy Dissipation Concept

The reaction rates $\mathscr{R}_j$ are modelled using the Eddy Dissipation Concept (EDC) [91], taking into account turbulence-chemistry interactions. The EDC model assumes that the reactions only occur in the fine scales, i.e. in the smallest turbulence scales. Their contribution to the mean reaction rate is accounted for through a pre-multiplying factor κ_{fs}, defined as:

$$\kappa_{fs}=\frac{\xi^2}{1-\xi^3} \quad (3.35)$$

In Eq. 3.35 the variable ξ represents the fine-structure region mass fraction, which is calculated as:

$$\xi=C_\xi\left(\frac{\nu\varepsilon}{\kappa^2}\right)^{\frac{1}{4}} \quad (3.36)$$

where C_ξ is a constant value equal to 2.1377 [35], ν is the kinematic viscosity, κ is the turbulent kinetic energy and ε is the rate of dissipation of turbulent kinetic energy. Moreover, the reactions are integrated over the time scale τ^*, which is defined as:

$$\tau^*=C_\tau\left(\frac{\nu}{\varepsilon}\right)^{\frac{1}{2}} \quad (3.37)$$

where C_τ is a constant value equal to 0.4082 [35].

Finally, the reaction rates $\mathscr{R}_j$ can be written as [6]:

$$\mathscr{R}_j = \bar{\rho}\,\kappa_{fs}\frac{Y_j^* - \tilde{Y}_j}{\tau^*} \tag{3.38}$$

where Y_j^* represents the fine-scale species mass fraction after reacting over the time scale τ^* in an adiabatic and isobaric perfectly stirred reactor, whose initial conditions are the species and temperature at the end of the previous time-step (or iteration, depending on whether the simulation is transient or steady).

Furthermore, the Ansys® Fluent implementation of the Eddy Dissipation Concept (EDC) differs from the reference formulation [53], as it does not account for the possibility that reactions may occur only within a fraction of the fine scales, among other modifications [35].

3.1.4 Large Eddy Simulation

A LES simulation is employed to provide high-fidelity reference data for the validation of the 2D axisymmetric RANS approximation in the mixing tube region of the AHEAD combustor, where experimental PIV data are not available. Unlike RANS, which models all turbulent scales, LES directly resolves the large energy-containing eddies while modelling only the small-scale turbulence effects through a subgrid-scale (SGS) model.

In the finite-volume framework of Ansys® Fluent, the filtering operation is implicitly provided by the computational grid, with the filter width determined by the local cell volume V [6]:

$$\bar{\beta}(x) = \frac{1}{V}\int_V \beta(x')\,dx' \tag{3.39}$$

where $\bar{\beta}$ denotes the filtered value of a generic scalar β. For variable-density flows, Favre-filtered quantities $\tilde{\beta} = \overline{\rho\beta}/\bar{\rho}$ are employed, consistent with the RANS formulation.

The filtered momentum equation introduces the subgrid-scale stress tensor τ_{ij}^{sgs}, whose deviatoric part is modelled using the Boussinesq hypothesis:

$$\tilde{\tau}_{ij}^{sgs} - \frac{1}{3}\tilde{\tau_{kk}}^{sgs}\delta_{ij} = -2\mu_{sgs}\left(\tilde{S}_{ij} - \frac{1}{3}\tilde{S}_{kk}\delta_{ij}\right) \tag{3.40}$$

where μ_{sgs} is the subgrid-scale viscosity. The Wall-Adapting Local Eddy-viscosity (WALE) model [108] is adopted for the SGS closure, providing appropriate near-wall behaviour without requiring explicit damping functions.

To reduce the computational cost associated with resolving the viscous sublayer, the Wall-Modeled LES (WMLES) approach with a near-wall RANS layer is employed [6]. This formulation blends the LES and RANS viscosities according to:

$$\nu_{blended} = f_{hy}\,\nu_{RANS} + \left(1 - f_{hy}\right)\nu_{LES} \tag{3.41}$$

where the blending function f_{hy} depends on the wall distance d_w and local grid spacing. In the near-wall RANS region, a mixing-length model with wall-damping provides the turbulent viscosity.

Turbulence-Chemistry Interaction: Partially Stirred Reactor

For the LES calculations, the turbulence-chemistry interaction is modelled using the Partially Stirred Reactor (PaSR) model [6], which offers improved accuracy over the standard EDC formulation for LES applications. The PaSR model [37] defines the reacting volume fraction κ_{PaSR} as:

$$\kappa_{\mathrm{PaSR}} = \frac{\tau_c}{\tau_c + \tau_{mix}} = \frac{1}{1 + Da} \tag{3.42}$$

where $Da = \tau_{mix}/\tau_c$ is the Damköhler number, τ_c is the chemistry time scale and τ_{mix} is the mixing time scale. The reaction time scale τ^* is computed as:

$$\tau^* = \min\left(\tau_c, \tau_{mix}\right) \tag{3.43}$$

The chemistry time scale τ_c is evaluated from the reaction rates of major species, while the mixing time scale τ_{mix} is computed from the geometric mean of the Kolmogorov time scale and the turbulence time scale [6, 84].

3.2 Numerical Investigation of Hydrogen Micromix Combustors

3.2.1 Motivation and Background

The micromix combustion principle aims for low-NO_x hydrogen combustion exploiting the mixing of multiple miniaturized fuel jets with combustor air to achieve short residence times in small diffusion flamelets [45]. This section presents a numerical investigation of converging nozzle configurations within a micromix combustor architecture, with the dual objectives of reducing pollutant formation and preventing flashback.

The fundamental challenge addressed in this investigation arises from the competing requirements of hydrogen combustor design. As established in Section 1.5.1, hydrogen's high laminar flame speed and molecular diffusivity create substantial flashback risk in premixed or partially premixed configurations. Simultaneously, the elevated adiabatic flame temperatures discussed in Section 2.2.1 promote thermal NO_x formation through the Zeldovich mechanism detailed in Section 2.3.1. The micromix approach offers a potential resolution to this dilemma by creating conditions that simultaneously reduce residence time in high-temperature zones (thereby limiting NO_x) and maintain high flow velocities to prevent upstream flame propagation.

Taking as a reference the simple injector configurations investigated by Khandelwal et al. [76], the use of a converging nozzle between the hydrogen injection points and the combustion zone has been investigated. The converging nozzle serves two primary functions: it accelerates the flow to decrease residence times in the combustion chamber, and it provides a controlled mixing region where partial premixing can occur before the mixture enters the combustion chamber. The simple hypothesis is that steeper nozzle angles will produce higher flow velocities and shorter residence times, thereby reducing both NO_x formation and flashback susceptibility.

This initial study served to understand the ability to validate results already present in the literature and, at the same time, to assess the possibility of evaluating NO_x through RANS simulations. The specific objectives of this investigation are: first, to quantify the effect of converging nozzle angle on NO and N_2O formation

across a range of equivalence ratios; second, to assess the flashback resistance provided by the nozzle acceleration.

3.2.2 Combustor Geometry and Configurations

The baseline combustor geometry follows the configuration presented by Khandelwal et al. [76], consisting of a 70 mm long combustor comprising a 20 mm cylindrical section and a 50 mm combustion zone. The geometry is shown in Fig. 3.1. Air enters through the cylindrical duct of 6.4 mm diameter, while hydrogen is injected through four equally spaced ducts of 0.6 mm diameter located at the mid-plane of the cylindrical section.

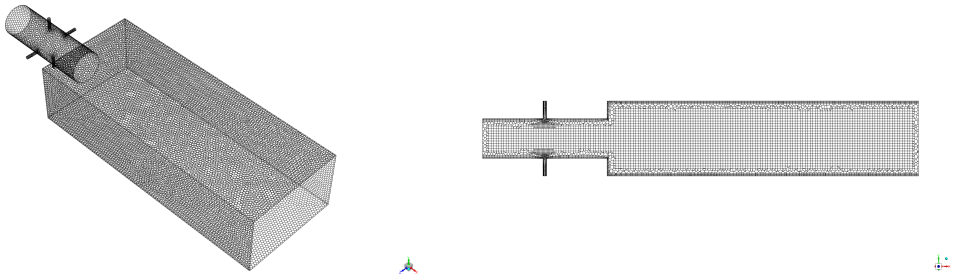

(a) Combustor CAD and surface mesh. (b) Poly-hexcore mesh on the X-Y plane.

Fig. 3.1 Fluid domain modelling for the geometry in Ref. [76].

The present study modifies this baseline by introducing a converging nozzle between the hydrogen injection points and the combustion zone, as illustrated schematically in Fig. 3.2. Three nozzle configurations are investigated, with half-angles of 10°, 15°, and 20°, while maintaining constant the distance between the injection points and the nozzle entrance as well as the nozzle length. The nozzle geometry accelerates the partially premixed hydrogen-air flow before it enters the combustion chamber, where flame is anchored at the nozzle exit.

3.2.3 Computational Methodology

Governing Equations and Turbulence Modelling

The simulations are performed in Ansys® Fluent 2022 R2 with the Reynolds-Averaged Navier-Stokes equations for steady-state, three-dimensional, compressible,

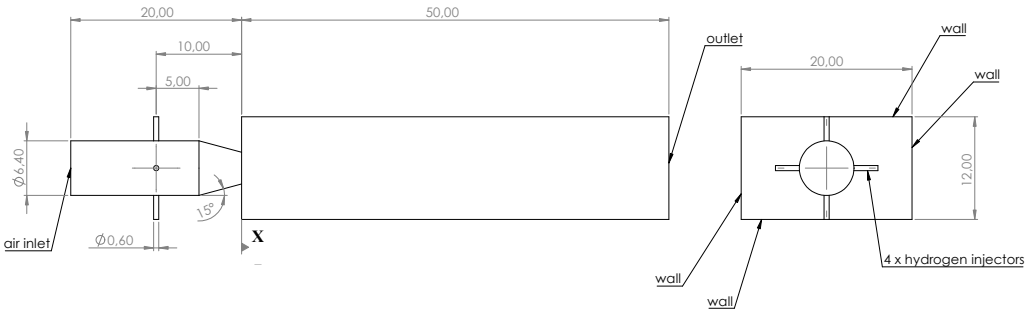

Fig. 3.2 Schematic of the micromix combustor geometry with converging nozzle. Dimensions shown for the 15° configuration. Boundary conditions: velocity inlet for air and hydrogen, pressure outlet at the combustion chamber exit, and adiabatic no-slip walls.

reacting flow, following the framework depicted in Sec. 3.1.1. Turbulence closure is provided by the k-ω Shear Stress Transport (SST) model [102].

The turbulent Schmidt and Prandtl numbers, which govern the turbulent transport of species and energy respectively, are set to their default values of 0.7 and 0.85. The sensitivity of hydrogen flame predictions to these parameters, especially to the Schmidt number, warrants consideration; however, for the present parametric study focused on relative comparisons between configurations, the default values provide an appropriate balance between accuracy and computational cost.

Turbulence-chemistry interactions are modelled using the Eddy Dissipation Concept (EDC) [91], shown in Sec. 3.1.3, and the chemical mechanism used is the one developed by Ó Conaire et al. [165], which comprises 10 species (H, H_2, O, O_2, OH, H_2O, N_2, HO_2, H_2O_2, Ar) and 19 reaction steps. This mechanism has been validated extensively against experimental data for laminar flame speeds, ignition delay times, and species concentrations across a wide range of conditions relevant to gas turbine combustors [163, 147, 109].

Pollutant Post-Processing

Following convergence of the main combustion simulation, pollutant formation is evaluated using a post-processing approach. This decoupled methodology is appropriate because NO_x species are present at trace concentrations that do not significantly affect the main flame structure or heat release. The transport equations for NO and N_2O, which account for convective and diffusive transport as well as

chemical source terms, are solved on the frozen flow field:

$$\frac{\partial}{\partial t}(\rho Y_{\mathrm{NO}}) + \nabla \cdot (\rho \vec{v} Y_{\mathrm{NO}}) = \nabla \cdot (\rho D_{\mathrm{m}} \nabla Y_{\mathrm{NO}}) + S_{\mathrm{NO}} \tag{3.44}$$

$$\frac{\partial}{\partial t}(\rho Y_{\mathrm{N_2O}}) + \nabla \cdot (\rho \vec{v} Y_{\mathrm{N_2O}}) = \nabla \cdot (\rho D_{\mathrm{m}} \nabla Y_{\mathrm{N_2O}}) + S_{\mathrm{N_2O}} \tag{3.45}$$

where ρ is the density, Y is the mass fraction, $\vec{v}$ is the velocity vector, and D_{m} is the mixture diffusivity.

Thermal NO_x Formation

Thermal nitric oxide formation is computed using the extended Zeldovich mechanism, presented in Section 2.3.1:

$$\mathrm{O} + \mathrm{N_2} \rightleftharpoons \mathrm{N} + \mathrm{NO} \tag{3.46}$$

$$\mathrm{N} + \mathrm{O_2} \rightleftharpoons \mathrm{O} + \mathrm{NO} \tag{3.47}$$

$$\mathrm{N} + \mathrm{OH} \rightleftharpoons \mathrm{H} + \mathrm{NO} \tag{3.48}$$

The rate of formation of NO_x via the thermal mechanism is significant only at high temperatures (above approximately 1800 K) due to the high activation energy required to break the strong N_2 triple bond. Applying a quasi-steady-state assumption for the intermediate nitrogen atom radical:

$$\frac{d[\mathrm{N}]}{dt} \approx 0 \tag{3.49}$$

the thermal NO formation rate becomes [6]:

$$\left(\frac{d[\mathrm{NO}]}{dt}\right)_{\text{thermal}} = 2k_{f,1}[\mathrm{O}][\mathrm{N_2}] \frac{1 - \dfrac{k_{r,1}k_{r,2}[\mathrm{NO}]^2}{k_{f,1}[\mathrm{N_2}]k_{f,2}[\mathrm{O_2}]}}{1 + \dfrac{k_{r,1}[\mathrm{NO}]}{k_{f,2}[\mathrm{O_2}] + k_{f,3}[\mathrm{OH}]}} \tag{3.50}$$

where concentrations are expressed in mol/m^3. The forward and reverse rate constants for reactions (3.46)–(3.48) are given by Arrhenius expressions [59]:

$$k_{f,1} = 1.8 \times 10^8 \exp(-38370/T) \qquad k_{r,1} = 3.8 \times 10^7 \exp(-425/T) \tag{3.51}$$

$$k_{f,2} = 1.8 \times 10^4 T \exp(-4680/T) \qquad k_{r,2} = 3.81 \times 10^3 T \exp(-20820/T) \tag{3.52}$$

$$k_{f,3} = 7.1 \times 10^7 \exp(-450/T) \qquad k_{r,3} = 1.7 \times 10^8 \exp(-24560/T) \tag{3.53}$$

where all rate constants have units of m^3/(mol·s) and T is the temperature in Kelvin.

N_2O-Intermediate Pathway

The N_2O intermediate pathway, which becomes significant at elevated pressures and moderate temperatures as discussed in Section 2.3.2, is included through the following mechanism [94]:

$$\mathrm{N_2 + O + M \rightleftharpoons N_2O + M} \tag{3.54}$$

$$\mathrm{N_2O + O \rightleftharpoons 2\,NO} \tag{3.55}$$

where M represents a third-body collision partner. Because reaction (3.54) involves third bodies, the mechanism is favoured at elevated pressures. Both reactions involve the oxygen radical O, making this mechanism important for oxygen-rich conditions. The contribution to NO formation from this pathway is:

$$\left(\frac{d[\mathrm{NO}]}{dt}\right)_{\mathrm{N_2O}} = 2\left(k'_{f,2}[\mathrm{N_2O}][\mathrm{O}] - k'_{r,2}[\mathrm{NO}]^2\right) \tag{3.56}$$

where $k'_{f,2}$ and $k'_{r,2}$ are the forward and reverse rate constants for reaction (3.55). The N_2O concentration is obtained by solving its transport equation (3.45), with the source term derived from reactions (3.54) and (3.55):

$$\frac{d[\mathrm{N_2O}]}{dt} = k'_{f,1}[\mathrm{N_2}][\mathrm{O}][\mathrm{M}] - k'_{r,1}[\mathrm{N_2O}][\mathrm{M}] - k'_{f,2}[\mathrm{N_2O}][\mathrm{O}] + k'_{r,2}[\mathrm{NO}]^2 \tag{3.57}$$

where the first two terms represent the forward and reverse contributions from reaction (3.54), and the last two terms correspond to the consumption and formation of N_2O via reaction (3.55). The third-body concentration [M] is computed as the total molar concentration of the mixture.

The forward and reverse rate constants for reactions (3.54) and (3.55) are given by the following Arrhenius expressions [6, 59]:

$$k'_{f,1} = 4.44 \times 10^{32} T^{-8.358} \exp(-28234/T) \tag{3.58}$$

$$k'_{r,1} = 4.00 \times 10^{7} \exp(-28234/T) \tag{3.59}$$

$$k'_{f,2} = 2.90 \times 10^{7} \exp(-11651/T) \tag{3.60}$$

$$k'_{r,2} = 1.45 \times 10^{-29} T^{9.259} \exp(-11651/T) \tag{3.61}$$

where T is the temperature in Kelvin. The rate constant $k'_{f,1}$ has units of $\mathrm{m^6/(mol^2 \cdot s)}$ due to the third-body dependence of reaction (3.54), while $k'_{f,2}$, $k'_{r,1}$, and $k'_{r,2}$ have units of $\mathrm{m^3/(mol \cdot s)}$. The appearance of [NO] in the reverse rate of reaction (3.55) introduces a coupling between the N_2O-intermediate pathway and the thermal NO_x mechanism.

Total NO Source Term

The total chemical source term for the NO transport equation (3.44) combines contributions from both pathways:

$$S_{\mathrm{NO}} = M_{w,\mathrm{NO}} \left[\left(\frac{d[\mathrm{NO}]}{dt} \right)_{\text{thermal}} + \left(\frac{d[\mathrm{NO}]}{dt} \right)_{\mathrm{N_2O}} \right] \tag{3.62}$$

where $M_{w,\mathrm{NO}} = 30.006$ kg/kmol is the molecular weight of NO. Similarly, the source term for the N_2O transport equation (3.45) is:

$$S_{\mathrm{N_2O}} = M_{w,\mathrm{N_2O}} \frac{d[\mathrm{N_2O}]}{dt} \tag{3.63}$$

where $M_{w,\mathrm{N_2O}} = 44.013$ kg/kmol.

3.2.4 Pollutant Turbulence-Chemistry Interaction

Turbulence-chemistry interactions for pollutant formation are modelled using a probability density function (PDF) approach. This treatment is essential because the NO_x formation rates exhibit a highly nonlinear dependence on temperature through the Arrhenius rate expressions, and using mean quantities alone would significantly

underpredict NO_x emissions. A joint PDF is assumed for temperature and OH mass fraction, as the thermal NO_x mechanism depends strongly on both variables. The two variables are treated as statistically independent, such that the joint PDF can be expressed as the product of the marginal distributions [6]:

$$P(T, Y_{OH}) = P_T(T) \cdot P_{Y_{OH}}(Y_{OH}) \tag{3.64}$$

Each marginal distribution is modelled using a beta-PDF, which is a two-moment function characterised by the local mean and variance of the respective variable [6]. Since the beta function is defined over the interval $[0,1]$, the temperature must be normalized using prescribed minimum and maximum bounds. The maximum temperature for normalization is determined locally as 1.1 times the mean temperature in each cell, ensuring adequate coverage of the temperature distribution while avoiding physically unrealistic integration bounds in cooler regions of the domain.

The mean source terms for the NO and N_2O transport equations are obtained by integrating the instantaneous source terms, defined in Equations (3.62) and (3.63), over the joint PDF:

$$\bar{S}_{NO} = \iint S_{NO}(T, Y_{OH})\, P(T, Y_{OH})\, dT\, dY_{OH} \tag{3.65}$$

$$\bar{S}_{N_2O} = \iint S_{N_2O}(T, Y_{OH})\, P(T, Y_{OH})\, dT\, dY_{OH} \tag{3.66}$$

where $S_{NO}(T, Y_{OH})$ and $S_{N_2O}(T, Y_{OH})$ are the instantaneous source terms evaluated at the instantaneous temperature T and OH mass fraction Y_{OH}. The integration is performed numerically using 30 discretization points for each variable.

The beta-PDF requires the local mean values and variances of the independent variables. However, the variance calculation differs between temperature and species mass fraction. The temperature variance σ_T^2 is computed by solving the following transport equation [6]:

$$\frac{\partial}{\partial t}\left(\rho \sigma_T^2\right) + \nabla \cdot \left(\rho \vec{v} \sigma_T^2\right) = \nabla \cdot \left(\frac{\mu_t}{\sigma_t} \nabla \sigma_T^2\right) + C_g \mu_t \left(\nabla \bar{T}\right)^2 - C_d \rho \frac{\varepsilon}{k} \sigma_T^2 \tag{3.67}$$

where $\bar{T}$ is the mean temperature, μ_t is the turbulent viscosity, k is the turbulent kinetic energy, and ε is its dissipation rate. The first term on the right-hand side represents turbulent diffusion of the variance, the second term accounts for variance

production due to mean temperature gradients, and the third term models the dissipation of temperature fluctuations at the small scales. The model constants take the values $\sigma_t = 0.85$, $C_g = 2.86$, and $C_d = 2.0$.

In contrast, the OH mass fraction variance $\sigma^2_{Y_{\mathrm{OH}}}$ is computed using an algebraic closure that assumes local equilibrium between production and dissipation of fluctuations [6]:

$$\sigma^2_{Y_{\mathrm{OH}}} = \frac{\mu_t}{\rho} \frac{k}{\varepsilon} \frac{C_g}{C_d} |\nabla \bar{Y}_{\mathrm{OH}}|^2 \tag{3.68}$$

This algebraic approximation avoids the computational cost of solving an additional transport equation and is the recommended approach when the PDF involves multiple independent variables. While the transported variance equation provides a more accurate representation of the temperature fluctuation field, the algebraic formulation offers a reasonable estimate of species fluctuations based on local gradients, which is particularly suitable for the post-processing treatment of NO_x prediction.

Numerical Implementation

All simulations are performed using Ansys® Fluent 2022 R2 using the pressure-based solver with the coupled algorithm for the pressure-velocity coupling [6]. The solution procedure begins with a non-reacting flow simulation to establish the velocity and mixing fields. Subsequently, a temperature of 1500 K is patched in the combustion zone downstream of the converging nozzle to initiate combustion. The simulation then proceeds until convergence, defined by mass imbalance less than 10^{-6} and stabilization of residuals.

Spatial discretization employs second-order upwind schemes [13] for all transported quantities, with gradients evaluated using the least-squares cell-based method [6]. The specific discretization schemes adopted for each quantity are summarized in Table 3.1.

Convergence is assumed to have been achieved when the mass imbalance is less than 10^{-6}, and the residuals oscillate around an equilibrium position or stabilize for steady-state simulations.

Table 3.1 Spatial discretization schemes used for CFD simulations.

Quantity	Spatial Discretization
Pressure	Coupled
Density	Second Order Upwind
Momentum	Second Order Upwind
Turbulent Kinetic Energy	Second Order Upwind
Specific Dissipation Rate	Second Order Upwind
Species	Second Order Upwind
Energy	Second Order Upwind

Mesh Size

The computational domain is discretised using poly-hexcore meshes generated within Ansys® Fluent Meshing, with total cell counts ranging from 3×10^5 to 2×10^6 depending on the configuration. To maintain computational efficiency while accurately resolving critical flow features, dynamic gradient-based mesh adaptation is executed every 10 iterations.

Refinement is triggered based on statistical thresholds: specifically, cells are refined where the gradient of the hydrogen mass fraction exceeds the domain mean plus 0.5 times the standard deviation ($\mu + 0.5\sigma$), ensuring accurate capture of upstream hydrogen mixing. Similarly, gradients of temperature and OH mass fraction are utilized to resolve the flame structure; notably, this criterion allows for both refinement and coarsening to dynamically optimize the grid density in reacting zones. While no maximum limit was imposed on the refinement levels, the minimum cell edge length was strictly constrained to 5×10^{-5} m to prevent prohibitive computational costs. This resolution is finer than the baseline hexahedral edge length of 2.5×10^{-4} m. These spatial resolution parameters are consistent with the grid sensitivity study reported by Khandelwal et al. [76].

3.2.5 Operating Conditions

The baseline operating conditions, summarized in Table 3.2, are selected to enable comparison with both the computational results of Khandelwal et al. [76] and the experimental measurements of Marek et al. [96] for lean direct injection (LDI) hydrogen combustors.

Table 3.2 Baseline operating conditions for micromix combustor simulations.

Parameter	Value
Air inlet velocity	30.5 m/s
Equivalence ratio	0.266, 0.300, 0.372
Operating pressure	6.9 bar
Air inlet total temperature	600 K
Hydrogen inlet total temperature	298.15 K
Converging nozzle half-angle	0° (baseline), 10°, 15°, 20°

The turbulence boundary conditions at the air inlet specify a hydraulic diameter of 6.4 mm and turbulence intensity of 10%, reflecting the elevated turbulence levels expected downstream of the compressor. At the hydrogen injection ports, a hydraulic diameter of 0.6 mm and turbulence intensity of 5% are specified.

The equivalence ratio of 0.300 corresponds to an experimental measurement point in Ref. [96], enabling direct validation. The other equivalence ratios (0.266 and 0.372) span a range representative of practical low-NO_x combustor operation.

3.2.6 Results: Effect of Converging Nozzle on Flame Structure

The introduction of a converging nozzle fundamentally alters the flame structure compared to the non-premixed baseline configuration. Figure 3.3 and Figure 3.4 present static temperature contours for the baseline (no nozzle) and 15° converging nozzle configurations at an equivalence ratio of 0.300.

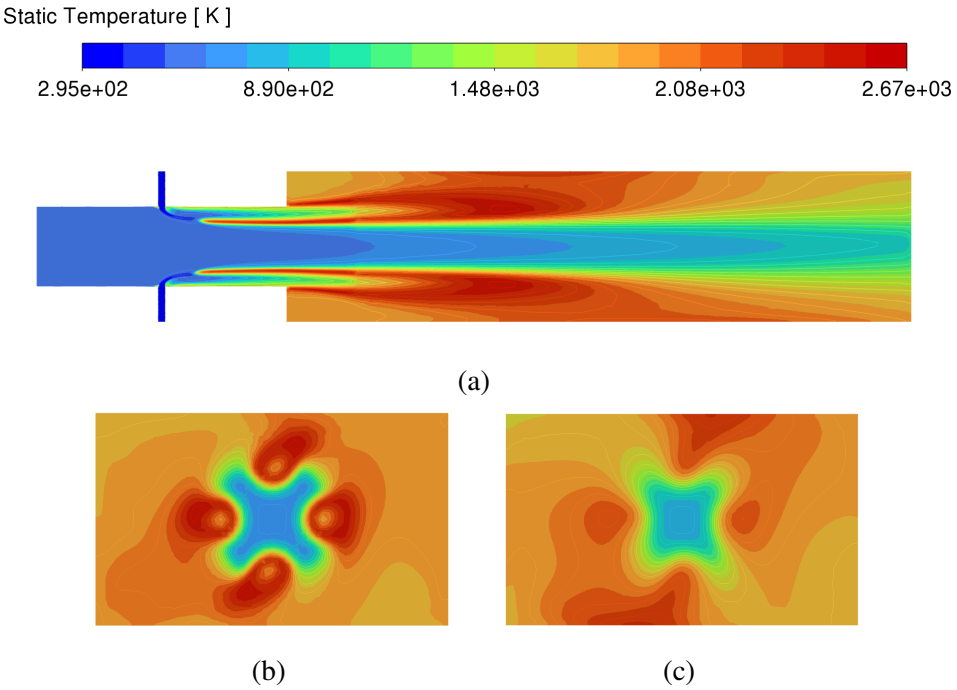

Fig. 3.3 Baseline configuration presented in Ref. [76], equivalence ratio = 0.300. Static Temperature contour plot: (a) along the combustor axis (X-Y plane), (b) X = 1 cm (Y-Z plane), (c) X = 2.54 cm (1 inch) (Y-Z plane).

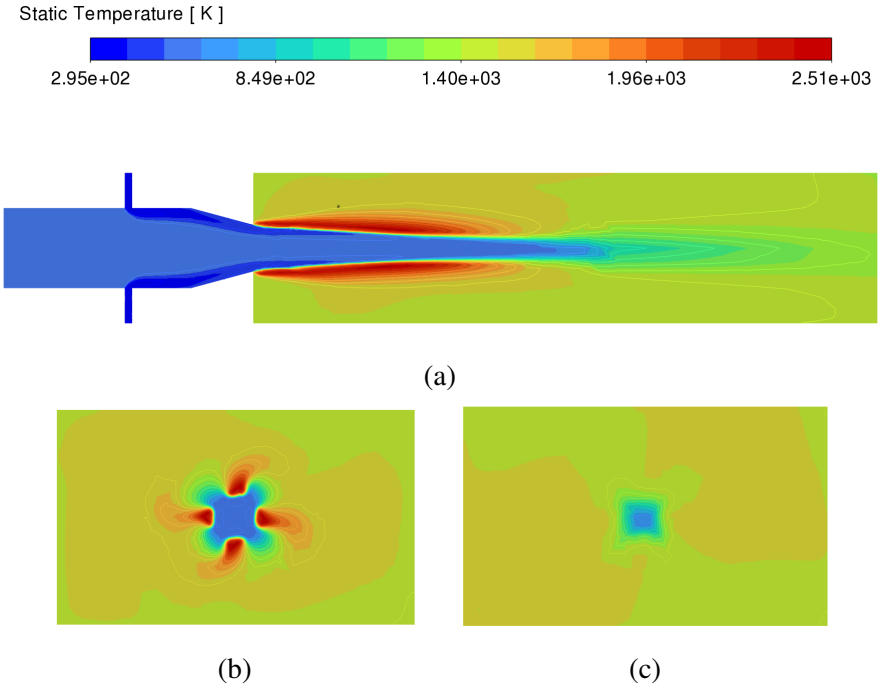

Fig. 3.4 Configuration with 15° converging nozzle, equivalence ratio = 0.300. Static Temperature contour plot: (a) along the combustor axis (X-Y plane), (b) X = 1 cm (Y-Z plane), (c) X = 2.54 cm (1 inch) (Y-Z plane).

Several qualitative differences are evident. The converging nozzle configuration produces a significantly shorter and thinner flame compared to the baseline. The partial premixing achieved within the nozzle enables more rapid combustion completion, while the flow acceleration reduces the residence time in high-temperature zones. These characteristics have direct implications for NO_x formation, as the thermal NO mechanism depends strongly on both peak temperature and residence time at elevated temperatures.

The radial temperature distributions at downstream stations ($X = 1$ cm and $X = 2.54$ cm) reveal that the converging nozzle configuration achieves more uniform temperature profiles, indicating better mixing and more complete combustion within a shorter distance. By the 1-inch station, the 15° nozzle configuration has reached a more homogeneous temperature field, whereas the baseline configuration still exhibits significant temperature gradients.

3.2.7 Results: Pollutant Emissions

Effect of Nozzle Angle and Equivalence Ratio

Table 3.3 presents the NO and N_2O concentrations (wet basis, mass-weighted average) at the 1-inch measurement plane for all configurations and equivalence ratios investigated. Results for the LDI combustor from Ref. [96] are included for comparison.

The results demonstrate a consistent trend: increasing the nozzle angle substantially reduces NO formation. The 20° configuration achieves NO concentrations of 3–6 ppm, representing a reduction of approximately 98% compared to the non-premixed baseline. This dramatic improvement results from the combined effects of flow acceleration (reducing residence time) and enhanced premixing (producing lower peak temperatures through more uniform combustion).

An inverse relationship is observed between NO and N_2O concentrations. As the converging nozzle angle increases and thermal NO is suppressed, the relative contribution of the N_2O intermediate pathway becomes more significant. This behaviour is consistent with the discussion in Section 2.3.2: the N_2O pathway is favoured at moderate temperatures (1200–1800 K) and elevated pressures, precisely the conditions achieved in the more aggressively premixed configurations. The N_2O

Table 3.3 NO and N_2O concentrations (ppm, wet basis) at $X = 2.54$ cm (1 inch) for different nozzle configurations and equivalence ratios.

Configuration	Equivalence ratio	NO (ppm)	N_2O (ppm)
Non-premixed (baseline)	0.266	171.5	0.7
	0.300	207.8	0.7
	0.372	271.3	0.9
10° converging nozzle	0.266	23.7	2.0
	0.300	42.1	1.4
	0.372	63.0	1.3
15° converging nozzle	0.266	10.5	3.5
	0.300	15.6	3.6
	0.372	17.8	3.7
20° converging nozzle	0.266	4.2	5.0
	0.300	3.3	4.0
	0.372	5.6	5.7
LDI [96] (wall measurement)	0.266	1.5	—
	0.300	2.1	—
	0.372	2.8	—

concentrations remain below 6 ppm in all cases, and total NO_x (NO + contribution from N_2O decomposition) remains well below the baseline non-premixed values.

Comparison with the LDI experimental data requires careful interpretation. The measurements in Ref. [96] were obtained near the wall using a gas analyzer, whereas the simulation results represent mass-weighted averages across the cross-section. Furthermore, the flame lengths differ between the micromix and LDI configurations. The 20° nozzle configuration achieves NO levels comparable to the LDI measurements, suggesting that the converging nozzle micromix concept can deliver competitive emissions performance.

3.2.8 Discussion

The results presented in this section demonstrate that converging nozzles offer a promising approach for achieving low-NO_x hydrogen combustion while maintaining flashback resistance. The 20° nozzle configuration achieves NO concentrations

comparable to experimental LDI injectors, with reductions of approximately 98% compared to non-premixed diffusion combustion.

The physical mechanisms underlying this improvement can be understood through the NO_x formation kinetics discussed in Section 2.3. The thermal NO formation rate depends exponentially on temperature and linearly on residence time at elevated temperatures. The converging nozzle reduces both factors: flow acceleration decreases residence time, while enhanced premixing produces more uniform temperature distributions with lower peak values. The result is a dramatic suppression of thermal NO.

The observed increase in N_2O with increasing nozzle angle indicates a shift in the dominant NO_x formation pathway. When thermal NO is suppressed, the pressure-dependent N_2O intermediate mechanism becomes the primary contributor. This observation has implications for mechanism selection in emissions modelling: accurate prediction of NO_x in low-emission hydrogen combustors requires inclusion of both thermal and N_2O pathways.

Several limitations of this investigation should be acknowledged:

- The simulations consider a single-injector configuration. In practical annular combustors, interactions between adjacent injector jets may affect both mixing and emissions.
- The RANS approach provides time-averaged solutions and cannot capture unsteady phenomena such as thermoacoustic oscillations or transient flashback events.
- Total pressure losses across the converging nozzle have not been quantified. The trade-off between emissions reduction and pressure loss must be considered in practical design.
- The comparison with LDI experimental data is complicated by differences in measurement methodology and flame structure.
- Flame anchoring at the injector plate can induce substantial thermal loads and generate surface hot spots, increasing the risk of ignition within the premixing section. Consequently, practical designs must incorporate effective cooling of the premixing section to mitigate this issue.

While acknowledging these limitations, the results establish converging nozzles as a potential design strategy for hydrogen micromix combustors to avoid flashback and reduce emissions.

3.2.9 Key Findings

The principal findings from this numerical investigation of hydrogen micromix combustors with converging nozzles are:

1. **NO reduction**: Increasing the converging nozzle angle from 0° to 20° reduces NO emissions by approximately 98%, from over 200 ppm to 3–6 ppm at representative conditions.
2. **N_2O pathway significance**: As thermal NO is suppressed, the N_2O intermediate pathway becomes the dominant NO_x formation mechanism. Total N_2O concentrations remain below 6 ppm.
3. **Flashback prevention**: The converging nozzle provides effective flashback resistance across all operating conditions investigated, with flames consistently stabilizing downstream of the nozzle exit.

These findings establish the foundation for further investigation of the micromix concept with converging nozzles. The demonstrated capability to achieve low NO_x emissions while preventing flashback addresses two of the primary challenges identified in Chapter 2 for hydrogen combustor development. The next section extends the investigation to swirl-stabilized configurations, where multi-fidelity modelling techniques are employed to efficiently explore the design space with limited experimental data.

3.3 Multi-Fidelity Modelling for Efficient Flashback Assessment

3.3.1 Motivation: From 3D to 2D Modelling with Data Correction

The micromix combustor investigation presented in Section 3.2 demonstrated that three-dimensional RANS simulations, while capable of capturing essential flow and combustion features, impose substantial computational demands. A single operating point required approximately $O(10^2\ \mathrm{CPU}\cdot\mathrm{h})$ on 30 cores, limiting the feasibility of extensive parametric studies across the design space.

This computational burden becomes particularly restrictive when investigating flashback limits, which require systematic exploration of operating conditions (equivalence ratio, air temperature, mass flow rate) to map the stability envelope. Experimental campaigns provide critical validation data but are themselves resource-intensive and necessarily limited in scope. The challenge, therefore, is to develop modelling strategies that maximize the value extracted from sparse experimental measurements while maintaining computational tractability for design exploration.

Multi-fidelity modelling offers a systematic framework for addressing this challenge by integrating information from sources of differing accuracy and cost [115]. The fundamental premise is that a correlation exists between low-fidelity predictions (computationally affordable but less accurate) and high-fidelity data (expensive but reliable), enabling construction of corrected models that approach high-fidelity accuracy at low-fidelity cost [44]. In combustion research, multi-fidelity approaches have been applied to instability characterization [159], emissions prediction [150], and combustor optimization [33].

This section presents a simplified multi-fidelity modelling approach, SimpleMF, developed by Optimad Srl [110] specifically for scenarios with severely limited high-fidelity training data. The methodology is demonstrated on the AHEAD lean premixed swirl-stabilized hydrogen burner [123], where axial flame position serves as an indicator of flashback susceptibility [124]. The approach employs two-dimensional axisymmetric RANS simulations as the low-fidelity model, a substantial simplifica-

tion from the three-dimensional treatment employed for the micromix combustor, combined with sparse experimental OH-PLIF measurements as high-fidelity data.

3.3.2 SimpleMF Model Formulation

The SimpleMF model assumes an approximately linear correlation between low- and high-fidelity outputs, augmented by a spatially-weighted correction field to capture residual discrepancies. This simplified structure reduces the number of hyperparameters compared to conventional multi-fidelity approaches [74, 116], improving training robustness when high-fidelity data are scarce.

Linear Correlation

Let x_{HF} and y_{HF} denote high-fidelity input and output data, and x_{LF}, y_{LF} the corresponding low-fidelity quantities. Assuming nested training sets $X_{HF} = \{x^i_{HF} : i = 1,\dots,n\} \subset X_{LF} = \{x^i_{LF} : i = 1,\dots,m\}$, the low-fidelity to high-fidelity relationship is approximated as:

$$y_{HF} \approx a \cdot y_{LF} + b \tag{3.69}$$

where $a, b \in \mathbb{R}$ are determined via least squares regression on the paired training data:

$$\begin{bmatrix} y^1_{LF} & 1 \\ \vdots & \vdots \\ y^n_{LF} & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} y^1_{HF} \\ \vdots \\ y^n_{HF} \end{bmatrix} \tag{3.70}$$

For any point x in the input domain, a first approximation $y^*_{HF}(x)$ is obtained as:

$$y^*_{HF}(x) = a \cdot y_{LF}(x) + b \tag{3.71}$$

where $y_{LF}(x)$ is computed directly from the low-fidelity model.

Radial Basis Function Correction

The linear approximation in Eq. (3.71) generally exhibits residual errors at the training points. These discrepancies $\delta^i = y^i_{HF} - y^*_{HF}(x^i_{HF})$ for $i = 1,\dots,n$ are interpolated using a radial basis function (RBF) with squared exponential kernel [24].

The correction field $\delta(x)$ provides an adjustment to the linear approximation:

$$y_{HF}(x) \approx y^*_{HF}(x) + \delta(x) \cdot w(x) \tag{3.72}$$

where the weight function $w(x)$ reduces the influence of the correction field at points distant from the training data:

$$w(x) = \frac{1}{1 + d_{\min}(x)} \quad , \quad d_{\min}(x) = \min_{i=1,\ldots,n} \|x - x^i_{HF}\|_2 \tag{3.73}$$

This weighting ensures that the approximation exactly matches the high-fidelity data at training points ($w = 1$ when $d_{\min} = 0$) while reducing the correction far from the training points. The SimpleMF model is implemented in the romBOX software [110].

3.3.3 AHEAD Burner Configuration

The AHEAD combustor, developed at TU Berlin [124], features a lean premixed swirl-stabilized configuration with axial hydrogen injection. The geometry comprises a radial swirl generator with adjustable blockage ratio (set to achieve geometric swirl number 0.9 for this study), a variable-diameter axial air inlet, and 16 radial hydrogen injection ports (diameter 0.8 mm) distributed around the fuel conduit. The mixing tube (diameter $D = 34$ mm) delivers the premixed reactants to a larger combustion chamber (diameter 105 mm). Additionally, 22 dilution air orifices (diameter 0.6 mm) at the mixing tube exit locally reduce the equivalence ratio near the wall, mitigating boundary layer flashback. Figure 3.5 presents the combustor geometry for the configuration investigated.

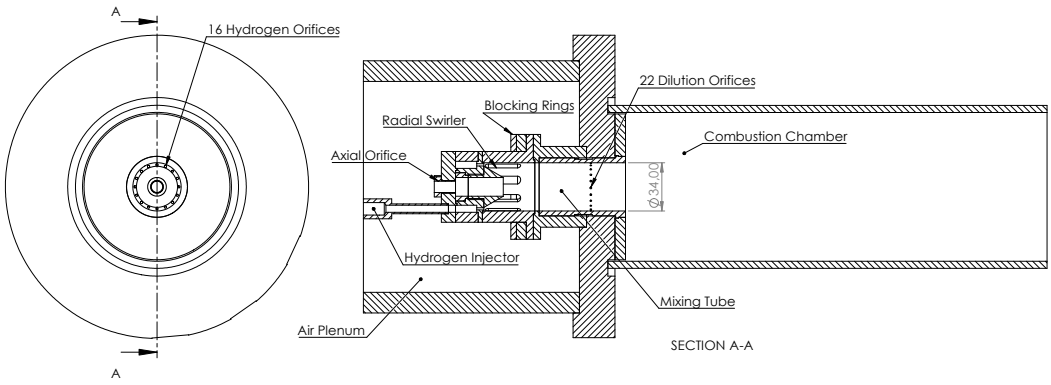

Fig. 3.5 CAD representation of the AHEAD combustor in the configuration with geometric swirl number 0.9 and axial orifice diameter 8 mm.

The axial injection ratio χ, defined as the ratio of axial air flow to total inlet air flow, is controlled by the variable axial orifice diameter and influences the velocity distribution at the mixing tube exit. All simulations employ $\chi = 12.5\%$, corresponding to an axial orifice diameter of 8.8 mm [125]. Hydrogen temperature at the injection ports is estimated from the air mass flow rate, equivalence ratio, and air temperature via an empirical correlation [124]:

$$T_{H_2} = \frac{a_{Fl}\, T_{\text{air}} + b_{Fl}\,(\dot{m}_{H_2}\, c_{p,H_2} - 0.5\, a_{Fl})}{0.5\, a_{Fl} + \dot{m}_{H_2}\, c_{p,H_2}} \tag{3.74}$$

with $a_{Fl} = 1.40$ W/°C, $b_{Fl} = 34.9\,^{\circ}$C, and c_{p,H_2} assumed constant over the relevant temperature range.

The operating parameters for this study cover the following ranges: inlet air temperature from 313 K to 693 K, inlet air mass flow rate from 80 kg/h to 295 kg/h, and equivalence ratios from 0.2 to 1. However, simulations at an equivalence ratio of 0.2 did not yield a stable flame. This limitation is likely associated with the intrinsic difficulties of RANS-based combustion modelling in reproducing the complex turbulence–chemistry interactions that dominate under ultra-lean conditions. At such low equivalence ratios, hydrogen flames exhibit pronounced thermo-diffusive and preferential-diffusion effects, strong sensitivity to strain and stretch, and a high susceptibility to local extinction and quenching. Therefore, operating conditions with an equivalence ratio of 0.2 are excluded from the present analysis.

The operating conditions considered here are derived from the flame distance data presented in Figures 9 and 10 of Ref. [124]. Figure 9 illustrates the flame distance

at a constant $u_0 = 70$ m/s for different inlet air temperatures and equivalence ratios, while Figure 10 depicts the flame distance at constant equivalence ratio ($\phi = 0.4$ and $\phi = 0.6$), with variations in inlet air temperature and bulk mixture velocity $u_0(\phi)$. It should be noted that some operating conditions appear in both figures but exhibit minor differences in the reported flame distances. While not explicitly stated in the original publication [124], these differences could stem from variations in the experimental configuration between the two test series (e.g., different mixing tube lengths or measurement locations) or from inherent measurement uncertainty. Despite these discrepancies, the fundamental flame behaviour captured in both datasets is considered to be sufficiently consistent and representative of the same burner concept, making them appropriate for training and validating the multi-fidelity model presented in this work. The parameter u_0 represents the mean velocity in the mixing tube when only air is present (i.e., equivalence ratio = 0) and is calculated as:

$$u_0 = \frac{\dot{m}_{\text{air}}}{A_{\text{mixing tube}}\,\rho_{\text{air}}} \tag{3.75}$$

The simulated operating conditions and corresponding experimental flame distances are summarized in Table 3.4. Here, x_{F}/D represents the flame distance from the exit of the mixing nozzle, normalized by the mixing nozzle diameter $D = 34\,\text{mm}$. All simulations assume an axial injection ratio $\chi = 12.5$, which, as described in Ref. [125], corresponds to an axial inlet diameter of 8.8 mm. These data represent the 53 high-fidelity operating conditions used in this study. Low-fidelity simulations under the same operating conditions are also computed to construct and validate the multi-fidelity model presented in Sec. 3.3.2.

3.3.4 Low-Fidelity Model: 2D Axisymmetric RANS

Geometric Simplification and Boundary Conditions

The three-dimensional AHEAD geometry is approximated as a two-dimensional axisymmetric domain, retaining axial and radial velocity components plus a tangential (swirl) velocity that varies with axial and radial position but not azimuthally:

$$\frac{\partial(\cdot)}{\partial\theta} = 0 \quad , \quad \vec{v}(x,r) = [v_x, v_r, v_\theta] \quad . \tag{3.76}$$

Table 3.4 Working conditions and corresponding experimental flame distances normalized with mixing tube diameter, extracted from Ref. [124].

ID	$\dot{m}_{\text{air}}$ (kg/h)	T_{air} (K)	ϕ	x_{F} / D	ID	$\dot{m}_{\text{air}}$ (kg/h)	T_{air} (K)	ϕ	x_{F} / D
1	255	313	0.3	0.277	28	80	453	0.4	0.171
2	255	313	0.4	0.285	29	130	453	0.4	0.293
3	255	313	0.5	0.351	30	150	453	0.4	0.299
4	255	313	0.6	0.450	31	205	453	0.4	0.353
5	255	313	0.7	0.627	32	230	453	0.4	0.417
6	180	453	0.3	0.247	33	93	623	0.4	0.302
7	180	453	0.4	0.262	34	110	623	0.4	0.299
8	180	453	0.5	0.275	35	148	623	0.4	0.350
9	180	453	0.6	0.361	36	82	693	0.4	0.306
10	180	453	0.8	0.615	37	100	693	0.4	0.317
11	180	453	1.0	0.991	38	133	693	0.4	0.324
12	130	623	0.3	0.229	39	130	313	0.6	0.389
13	130	623	0.4	0.233	40	184	313	0.6	0.479
14	130	623	0.5	0.216	41	220	313	0.6	0.529
15	130	623	0.6	0.250	42	295	313	0.6	0.558
16	130	623	0.8	0.474	43	80	453	0.6	0.293
17	130	623	1.0	0.866	44	130	453	0.6	0.399
18	116	693	0.3	0.241	45	150	453	0.6	0.389
19	116	693	0.4	0.228	46	205	453	0.6	0.442
20	116	693	0.5	0.220	47	230	453	0.6	0.555
21	116	693	0.6	0.251	48	93	623	0.6	0.349
22	116	693	0.8	0.463	49	110	623	0.6	0.342
23	116	693	1.0	0.807	50	148	623	0.6	0.362
24	130	313	0.4	0.257	51	82	693	0.6	0.339
25	184	313	0.4	0.313	52	100	693	0.6	0.346
26	220	313	0.4	0.342	53	133	693	0.6	0.346
27	295	313	0.4	0.375					

This approximation reduces computational cost compared to full three-dimensional treatment while preserving the essential swirling flow physics. The geometry is constructed such that inlet areas in the 2D representation match the total flow area of the actual 3D combustor, ensuring equivalent inlet velocities. Figure 3.6 illustrates the 2D computational domain with dimensions and boundary conditions.

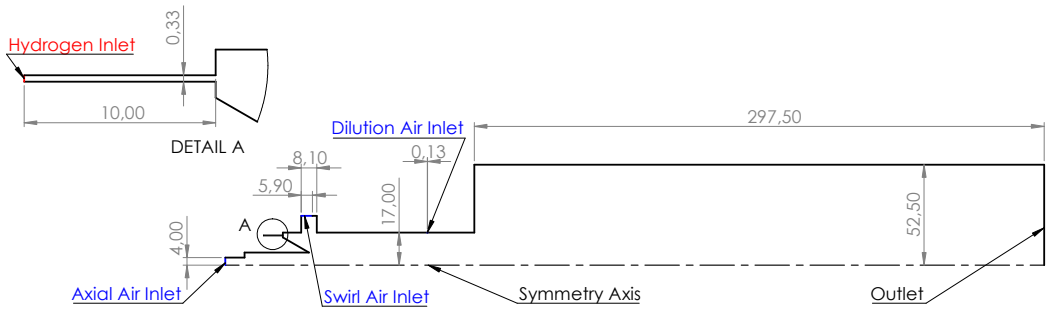

Fig. 3.6 Two-dimensional axisymmetric computational domain for the AHEAD combustor, showing dimensions and boundary conditions. The radial swirler, hydrogen injection, and dilution air are represented as equivalent area inlets

The radial swirl generator is modelled via a tangential velocity inlet followed by a wall to represent blockage. Since the continuous annular surface in 2D represents the distributed swirler slots in 3D, the inlet length in the simulation differs from the actual slot length; the geometric centroid of the actual slot is retained in the 2D approximation. The same centroid-preserving approach is applied to the hydrogen injection and dilution air inlets.

Mass flow rates at the boundaries are determined from the axial injection ratio definition:

$$\begin{cases} \dot{m}_{\text{dil}} = 0.025\,\dot{m}_{\text{tot}} \\ \dot{m}_{\text{swirl}} = (1-\chi)\,(\dot{m}_{\text{tot}} - \dot{m}_{\text{dil}}) \\ \dot{m}_{\text{axial}} = \dot{m}_{\text{tot}} - \dot{m}_{\text{swirl}} - \dot{m}_{\text{dil}} \end{cases} \tag{3.77}$$

where the 2.5% dilution air fraction is inferred from the ratio of dilution orifice area to total inlet area.

Simulation Framework

The RANS simulations are performed in Ansys® Fluent 2024 R1 following the 2D framework depicted in Sec. 3.1.2. Turbulence is modelled using the κ-ω SST

formulation [102], while turbulence-chemistry interaction is treated via the Eddy Dissipation Concept [91], following the approach employed for the micromix combustor in Section 3.2. Also in this case, hydrogen oxidation chemistry employs the Ó Conaire mechanism [165] (10 species, 19 reactions), validated extensively for hydrogen combustion [163, 147, 109].

The tangential velocity boundary condition at the swirl inlet requires calibration, as the geometric swirl number cannot be directly applied to the modified 2D geometry. A parametric study varying swirl ratio (0.75–0.85) and turbulent Schmidt number (0.1–0.7) via Latin Hypercube Sampling on 100 simulations identified optimal values of 0.8 and 0.7, respectively, based on minimizing L^2 error against experimental PIV velocity fields. These spatially uniform parameters represent a compromise: lower Schmidt numbers improve flame position prediction but may degrade mixing predictions in fuel-lean regions, while higher values better reproduce velocity fields globally. The calibrated parameters exhibit some Reynolds number dependence, potentially limiting accuracy at operating conditions significantly different from the calibration point, consistent with known swirler performance characteristics [130].

Mesh and Numerical Methods

The computational mesh employs approximately 2×10^5 poly-hexcore cells with three refinement zones: finest resolution in the hydrogen injection tube (mean cell size 3.42×10^{-10} m^2), intermediate refinement through the air inlet and flame zone (2.91×10^{-8} m^2), and coarsest cells downstream where gradients diminish (1.77×10^{-7} m^2). The mean y^+ across all walls is 0.51, with maximum values below 10. Grid independence was verified by refining each cell into four elements; velocity profiles and flame positions differ negligibly between baseline and refined meshes, confirming adequate resolution. Figure 3.7 illustrates the mesh topology and refinement strategy.

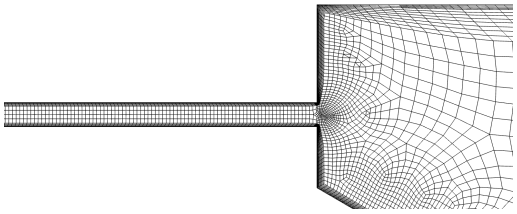

(a) Refined mesh around the hydrogen injection tube.

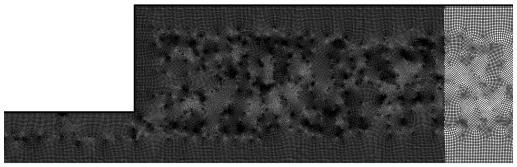

(b) Difference in mesh size in the combustion chamber.

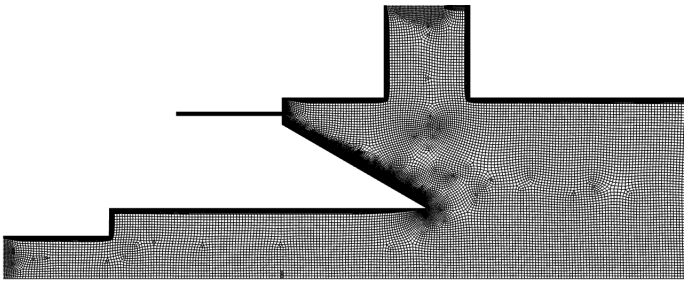

(c) Mesh topology at the air inlet and hydrogen injection region.

Fig. 3.7 Computational mesh features illustrating the baseline mesh.

Simulations are performed in Ansys® Fluent 2024 R1 using the pressure-based coupled solver with second-order upwind spatial discretization and least-squares cell-based gradient evaluation [6], as in the previous test case. Turbulent Prandtl and Schmidt numbers are set to 0.85 and 0.7, respectively, the latter informed by the calibration study. Average computational time per operating point is $O(7\ \text{CPU} \cdot \text{h})$ on 30 Intel Xeon Gold 6338 cores.

3.3.5 Low-Fidelity Model Validation

Velocity Field Comparison

Figure 3.8 compares axial velocity contours between experimental PIV measurements and 2D RANS predictions at $\dot{m}_{\text{air}} = 130$ kg/h, $T_{\text{air}} = 623$ K, $\phi = 0.6$. Despite the geometric simplifications, the CFD captures the essential flow structure: inner and outer recirculation zones (delineated by zero-velocity contours), and radial flow spreading towards the combustor walls. Two discrepancies are evident: the CFD positions the inner recirculation zone slightly upstream of the experimental

measurement, and predicts reduced radial spreading of the jet exiting the mixing tube.

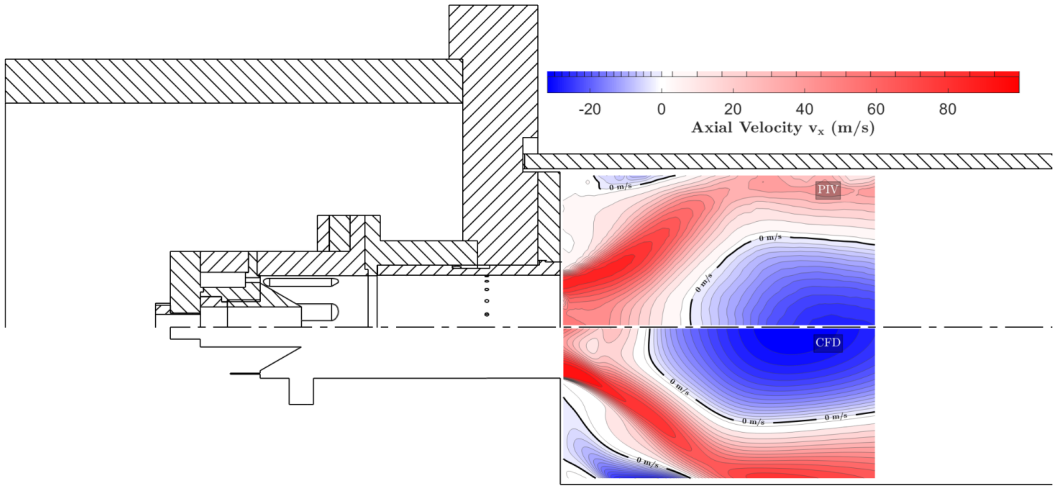

Fig. 3.8 Axial velocity contours from PIV (upper) and 2D RANS (lower). Thick black lines indicate zero axial velocity. Near-wall region omitted due to PIV measurement constraints.

Quantitative velocity profiles at multiple axial stations (Figs. 3.9, 3.10) confirm that the RANS model reproduces velocity magnitudes and spatial distributions reasonably well, though with systematic offsets in recirculation zone position. At the combustion chamber entrance ($x/D = 0.05$), predicted axial velocities slightly overpredict peak values while radial velocities match closely. Further downstream, the zero-velocity crossing in the radial profiles occurs at similar radial positions in both PIV and CFD, indicating comparable radial extent of the recirculation zones despite axial displacement. Mesh refinement produces negligible change in velocity predictions, confirming grid adequacy.

Flame Structure and Position

Figure 3.11 presents heat release contours for varying equivalence ratio at constant air conditions ($T_{\text{air}} = 453$ K, $\dot{m}_{\text{air}} = 180$ kg/h). Each contour is independently normalized by the maximum heat release rate of the respective case, so that the colour scale highlights flame structure rather than absolute magnitude. The simulations reproduce the experimentally observed trend [124]: the flame lifts off at lean conditions and progressively anchors closer to the mixing tube exit as equivalence ratio increases.

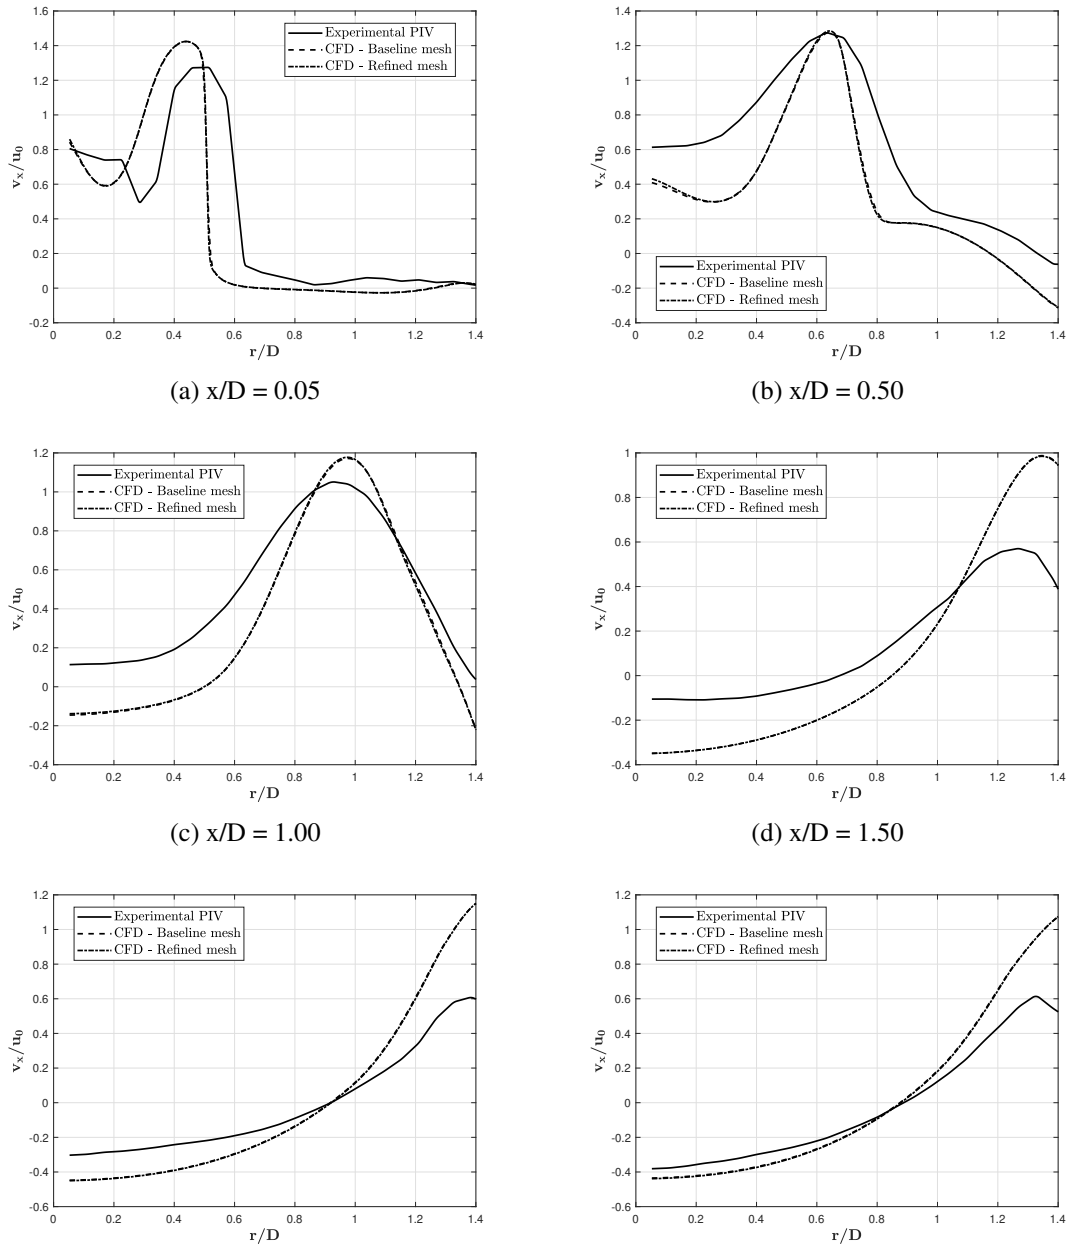

(a) x/D = 0.05

(b) x/D = 0.50

(c) x/D = 1.00

(d) x/D = 1.50

(e) x/D = 2.00

(f) x/D = 2.50

Fig. 3.9 Comparison of axial velocity v_x, normalized by velocity u_0, between PIV and CFD data at different axial locations x and radius r, normalized by mixing tube diameter D.

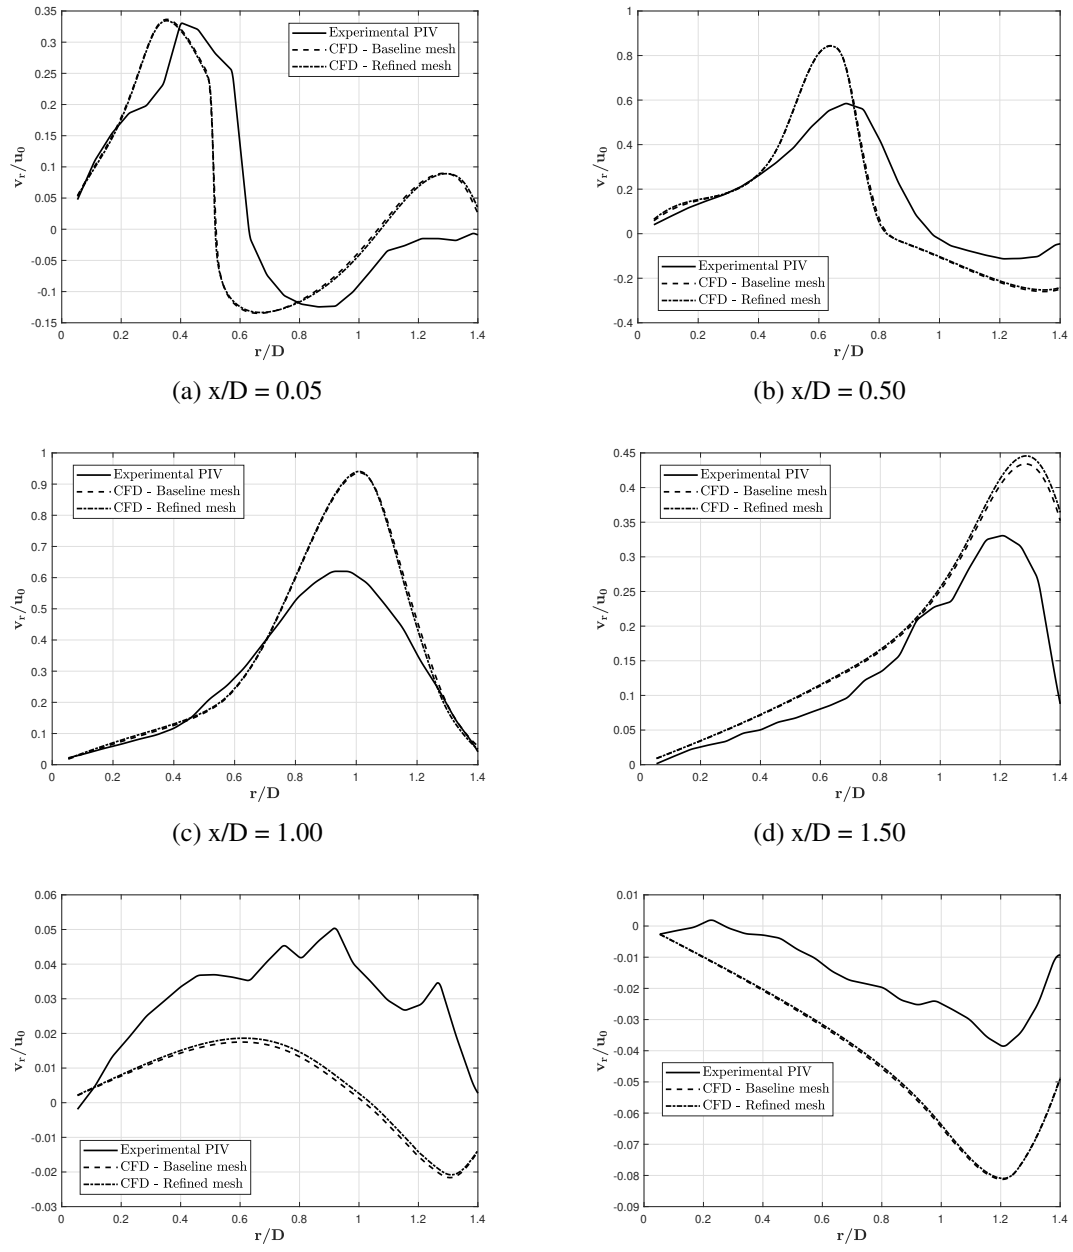

(a) x/D = 0.05

(b) x/D = 0.50

(c) x/D = 1.00

(d) x/D = 1.50

(e) x/D = 2.00

(f) x/D = 2.50

Fig. 3.10 Comparison of radial velocity v_r, normalized by velocity u_0, between PIV and CFD data at different axial locations x and radius r, normalized by mixing tube diameter D.

Experimentally, flame attachment occurs at $\phi = 0.5$; the RANS model predicts attachment at $\phi = 0.6$. At stoichiometric conditions, the RANS flame exhibits a flatter surface compared to the pronounced curvature observed experimentally, likely reflecting limitations of the 2D axisymmetric approximation and RANS turbulence closure in capturing three-dimensional flame-vortex interactions.

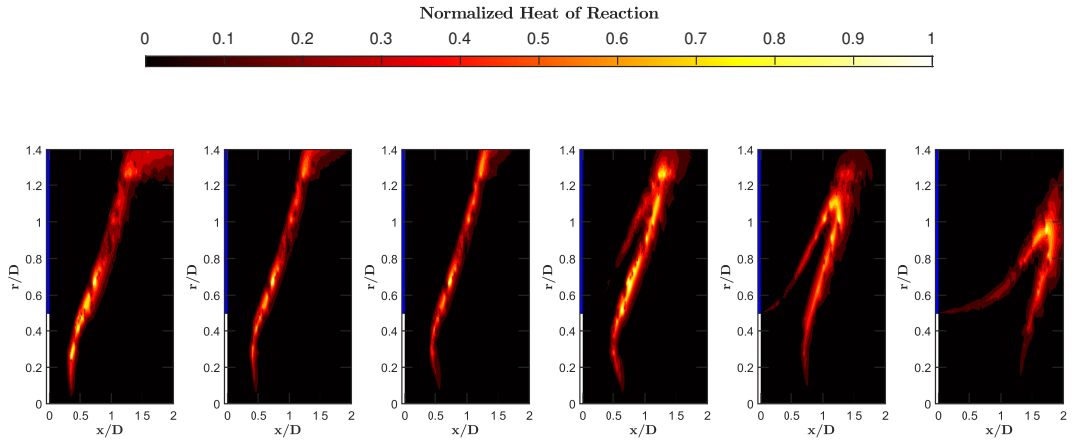

(a) $\phi = 0.3$ (b) $\phi = 0.4$ (c) $\phi = 0.5$ (d) $\phi = 0.6$ (e) $\phi = 0.8$ (f) $\phi = 1.0$

Fig. 3.11 CFD flame visualization through the normalized heat of reaction contours for six equivalence ratios, $\phi = 0.3$–1.0. Simulations performed at $T_{\text{in}} = 453\,\text{K}$ and $\dot{m}_{\text{air}} = 180\,\text{kg/h}$. The axial station $x/D = 0$ represents the exit of the mixing tube and the beginning of the combustion chamber. The blue line indicates the combustion chamber wall, while the white line follows the mixing tube radius.

Flame position is quantified as the axial location of maximum OH mass fraction gradient, normalized by mixing tube diameter. Figure 3.12(a) compares the complete low-fidelity (2D RANS) and high-fidelity (experimental) datasets across all 53 operating conditions. The CFD reproduces qualitative trends, flame distance increasing with equivalence ratio at constant temperature, but exhibits systematic offset. The Pearson correlation coefficient of 0.77 (Fig. 3.12b) confirms the assumption of approximate linear correlation underlying the SimpleMF model.

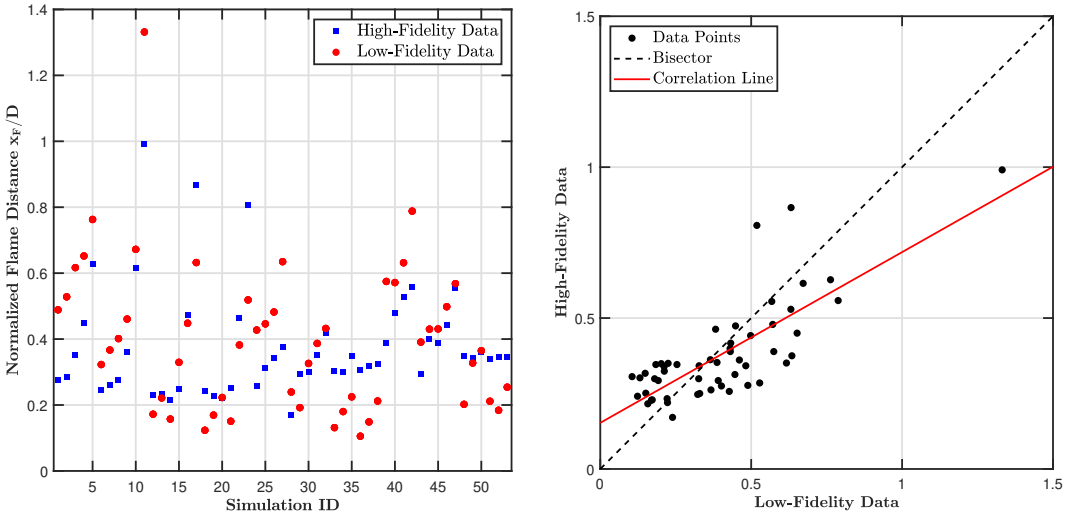

(a) Direct comparison

(b) Linear correlation

Fig. 3.12 Comparison of low-fidelity (2D RANS) and high-fidelity (experimental) flame positions across all operating conditions, demonstrating the approximate linear relationship

3.3.6 Comparison against LES

The experimental PIV data presented in Section 3.3.5 provides velocity measurements exclusively within the combustion chamber, leaving the flow structures inside the swirler and mixing tube unvalidated. This region is critical for the 2D axisymmetric simplification, as the swirl generation mechanism is inherently three-dimensional. To assess the fidelity of the 2D RANS approximation in these inaccessible regions, qualitative comparisons are performed against a three-dimensional Large Eddy Simulation (LES) conducted at the following operating conditions: $\dot{m}_{\text{air}} = 180$ kg/h, $T_{\text{air}} = 453$ K, $\phi = 0.6$.

The LES is performed in Ansys® Fluent 2024 R1 and employs the WALE (Wall-Adapting Local Eddy-viscosity) subgrid-scale model [108] combined with wall-modeled LES (WMLES) treatment [6], using a computational mesh of approximately 37 million cells. The mesh is spatially refined to resolve the relevant flow features: a characteristic cell size of 0.25 mm is adopted in the vicinity of the H_2 injectors, 0.50 mm within the swirler passages, the mixing tube and in the first section of the combustion chamber where the flame is anchored, 1.00 mm in the intermediate region of the combustion chamber, and 2.50 mm in the far downstream portion. The simulations are performed with a CFL number equal to 15 and mean quantities

are evaluated over a physical time of 0.5 ms. Spatial discretization employs the Bounded Central Differencing (BCD) scheme, which provides second-order accuracy while maintaining boundedness of transported scalars. Complete details of the LES methodology are provided in Section 3.1.4.

Velocity Field

Figure 3.13 compares velocity magnitude contours between LES (instantaneous and time-averaged) and 2D RANS simulations on a plane passing through the combustor axis.

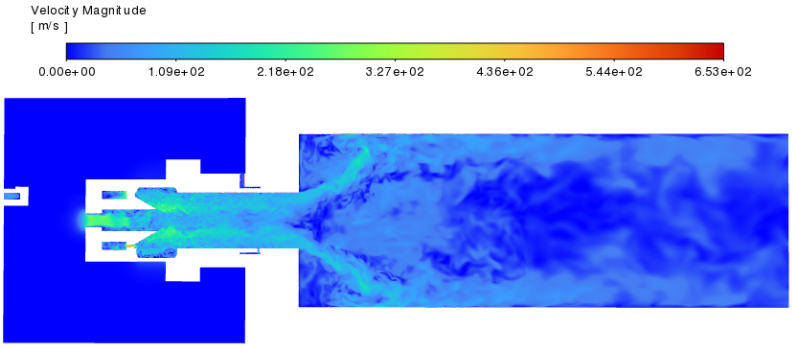

(a) LES instantaneous velocity magnitude.

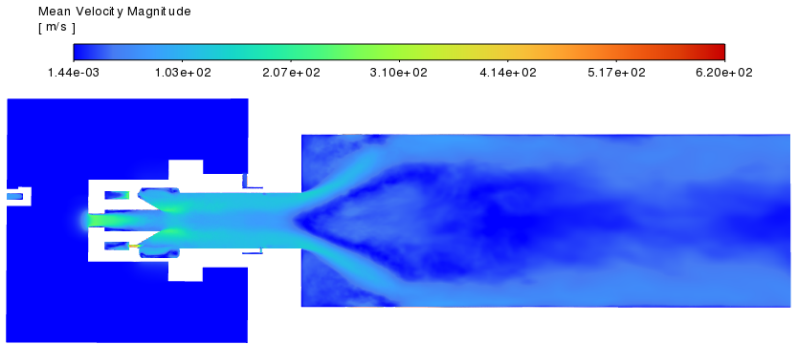

(b) LES time-averaged velocity magnitude.

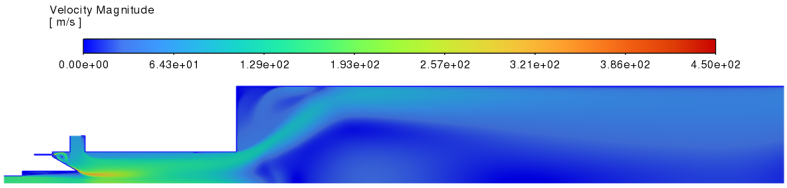

(c) 2D RANS velocity magnitude.

Fig. 3.13 Comparison of velocity magnitude contours between LES and 2D RANS.

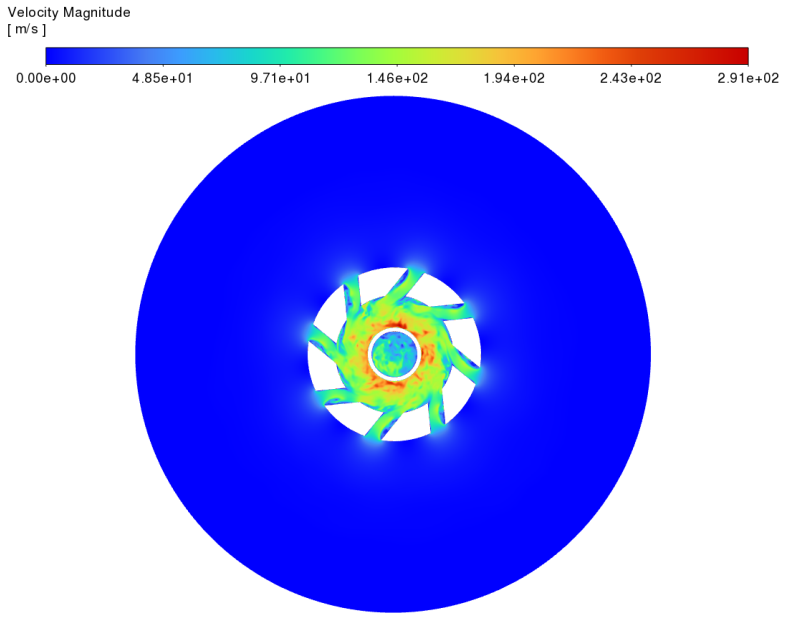

Fig. 3.14 Section of the radial swirler in the LES simulation.

The instantaneous LES field (Fig. 3.13a) reveals the turbulent nature of the swirling flow, with clearly visible coherent structures and velocity fluctuations throughout the mixing tube and combustion chamber. The time-averaged LES velocity field (Fig. 3.13b) exhibits a flow structure that is qualitatively similar to the 2D RANS prediction (Fig. 3.13c). Both approaches capture the high-velocity jet emanating from the mixing tube exit and its subsequent expansion into the combustion chamber. The flow acceleration through the annular passages surrounding the hydrogen injection tube is reproduced in both simulations. In the combustion chamber, both models predict a central jet region surrounded by lower velocity recirculation zones, consistent with the swirl-induced flow pattern.

Notable differences emerge in the swirler region. The LES resolves the discrete passages of the radial swirl generator, while the 2D RANS represents these as an equivalent continuous annular inlet. This simplification obviously affects the predicted tangential velocity distribution entering the mixing tube; additionally, the LES simulation shows separation in the swirler as depicted in Fig. 3.14. However, the mixing tube flow shows reasonable agreement between the time-averaged LES and RANS, suggesting that the 2D approximation adequately captures the mean flow development despite the geometric simplification.

Hydrogen Distribution

The hydrogen mixing process is compared in Figure 3.15, which displays hydrogen mass fraction contours for both modelling approaches. The instantaneous LES field (Fig. 3.15a) reveals highly wrinkled iso-concentration surfaces, indicating the turbulent mixing between hydrogen and air within the mixing tube. Hydrogen is injected through multiple tubes visible as discrete high-concentration regions at the injection plane.

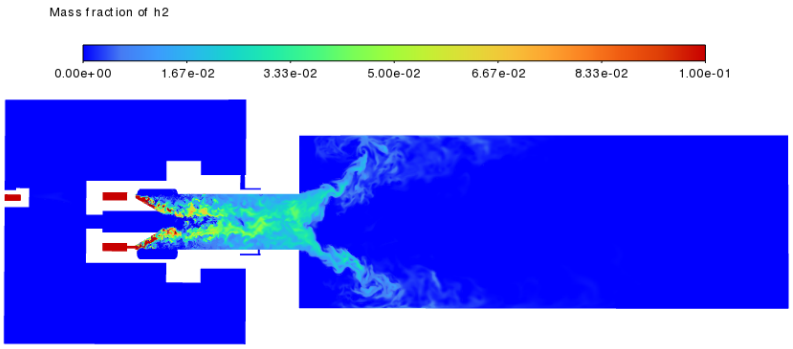

(a) LES instantaneous hydrogen mass fraction.

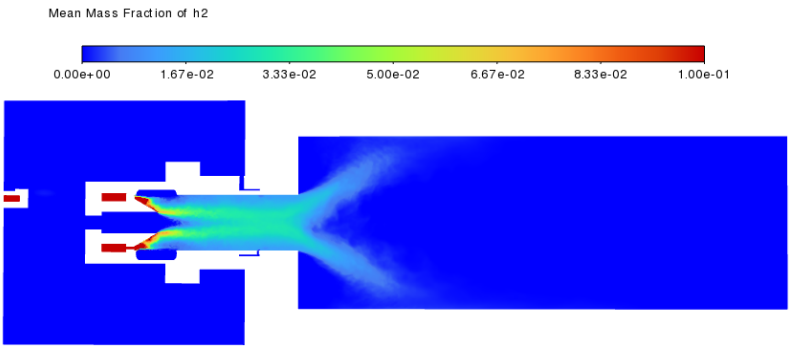

(b) LES time-averaged hydrogen mass fraction.

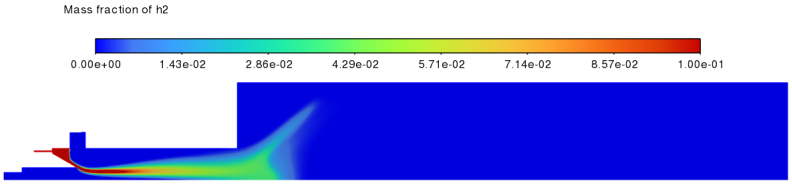

(c) 2D RANS hydrogen mass fraction.

Fig. 3.15 Comparison of hydrogen mass fraction distribution between LES and 2D RANS.

The time-averaged LES hydrogen distribution (Fig. 3.15b) shows a more diffuse pattern than the instantaneous snapshot, with hydrogen concentration gradually decreasing along the mixing tube as mixing with air progresses. The 2D RANS prediction (Fig. 3.15c) captures the overall mixing behaviour qualitatively, with hydrogen concentration highest near the injection plane and then a higher concentration of hydrogen mass fraction along the mixing tube axis. However, the RANS model predicts a narrower hydrogen plume compared to the LES, consistent with the lower turbulent diffusion typically observed in RANS simulations of swirling flows.

Both models indicate that hydrogen is largely consumed before reaching the combustion chamber outlet, confirming complete combustion under the simulated operating conditions. The region of unburned hydrogen extends slightly further into the combustion chamber in the LES results, suggesting that the RANS model may overpredict the reaction rates in the flame zone.

Temperature Field

Figure 3.16 presents temperature contours comparing the thermal field predictions. The instantaneous LES temperature field (Fig. 3.16a) displays the complex flame structure characteristic of turbulent premixed combustion, with a highly corrugated flame front separating the cold premixed reactants from the hot combustion products. The flame stabilizes at the mixing tube exit, where the flow expansion and recirculation zones provide aerodynamic anchoring.

The time-averaged LES temperature field (Fig. 3.16b) reveals a more gradual transition from cold reactants to hot products compared to the instantaneous field, reflecting the spatial averaging of the fluctuating flame position. The 2D RANS prediction (Fig. 3.16c) shows a similar overall pattern, with the premixed fuel-air mixture remaining at the inlet temperature through the mixing tube and temperature rising sharply at the flame front. The flame position, indicated by the steep temperature gradient, is located slightly downstream of the mixing tube exit in both the time-averaged LES and RANS solutions, with the RANS being further downstream. Indeed, as shown in Fig. 3.12, 2D RANS overestimates flame position in this operating condition.

A notable observation concerns the peak temperature. The LES predicts instantaneous temperatures exceeding 2100 K locally (Fig. 3.16a), consistent with

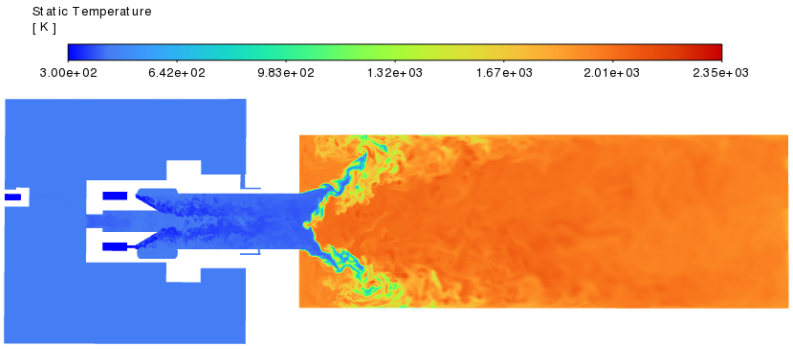

(a) LES instantaneous static temperature.

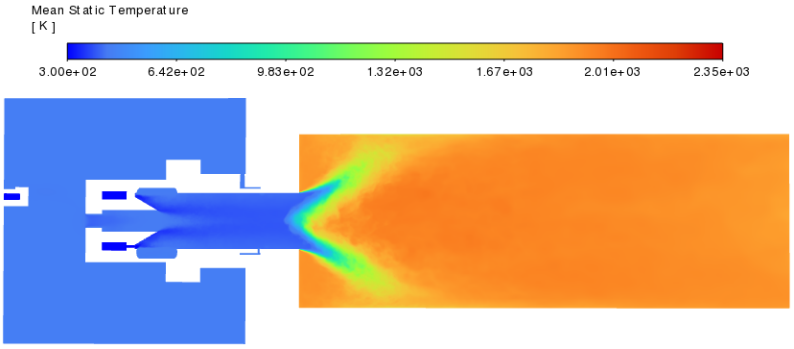

(b) LES time-averaged static temperature.

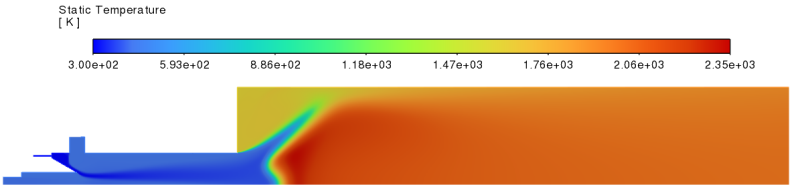

(c) 2D RANS static temperature.

Fig. 3.16 Comparison of static temperature distribution between LES and 2D RANS.

super-adiabatic burning in wrinkled flame elements. The time-averaged LES peak temperature is lower due to averaging over the flame surface fluctuations. The RANS simulation predicts a higher peak temperature and a colder outer recirculation zone: this difference seems to arise from the reduced diffusion of hydrogen in the mixing tube, leading to a higher equivalence ratio in the flame zone. Both approaches predict near-complete temperature uniformity in the downstream portion of the combustion chamber, indicating thorough mixing of combustion products.

Summary of LES-RANS Comparison

The qualitative comparison between three-dimensional LES and 2D axisymmetric RANS demonstrates that the simplified RANS approach captures the essential flow and flame characteristics for the AHEAD combustor under the investigated operating conditions. The principal findings are:

1. The time-averaged velocity field from LES exhibits qualitative agreement with the 2D RANS solution, also in the mixing tube and flame zone where PIV validation data are unavailable. The flow acceleration through the swirl passages and jet expansion at the combustor entrance are reproduced.

2. Hydrogen mixing patterns show reasonable correspondence between the models, though the RANS predicts a narrower fuel plume. These differences are attributable to the limitations of gradient-diffusion turbulent transport models in swirling flows for 2D RANS.

3. The flame position and temperature field structure are consistent between LES and RANS, supporting the use of 2D RANS for flame position prediction in the multi-fidelity framework. The RANS model captures the essential flame anchoring mechanism at the mixing tube exit.

The favourable agreement supports the validity of the 2D RANS approach as a computationally efficient low-fidelity model for the multi-fidelity framework, recognising that systematic biases exist and are corrected through the SimpleMF calibration procedure described in Section 3.3.7. It should be noted that the present comparison is qualitative and is conducted against a single LES performed by the author at one operating condition. A more quantitative verification is carried out in Chapter 4, where axial and radial velocity profiles from both RANS–2D and RANS–3D EARSM are compared against PIV data and an LES reference from the literature [5] at five axial stations (Section 4.3.6). The results of that comparison confirm that the discrepancy in the velocity observed near the combustion chamber entrance is not specific to the 2D approximation, pointing to limitations in reproducing the experimental boundary conditions via numerical analysis.

3.3.7 Multi-Fidelity Model Performance

To evaluate SimpleMF predictive accuracy, 15 randomly selected operating points serve as validation data (IDs: 7, 8, 10, 13–16, 29, 32, 33, 44–48), chosen to avoid parameter space boundaries. The remaining 38 points constitute potential training sets, constructed iteratively by adding the point furthest from existing training points in normalized parameter space. This procedure is repeated with different initial points to generate error statistics.

SimpleMF is benchmarked against three alternatives: single-fidelity Kriging [122], and two multi-fidelity methods: CoKriging [74, 54] and Nonlinear AutoRegressive Gaussian Process (NARGP) [116]. All models employ squared exponential kernels with hyperparameters determined via maximum likelihood estimation using BFGS optimization with five random restarts.

Figure 3.17 presents root mean square error (RMSE) distributions as training set size increases. SimpleMF consistently achieves lower median error than both Kriging fitted to high-fidelity data alone and uncorrected low-fidelity predictions. Remarkably, SimpleMF remains effective with training sets as small as three points, comparable to the number of input parameters (air mass flow, air temperature, equivalence ratio). CoKriging and NARGP, despite their greater flexibility, exhibit inferior performance at small training set sizes, attributed to difficulties in reliably estimating their numerous hyperparameters.

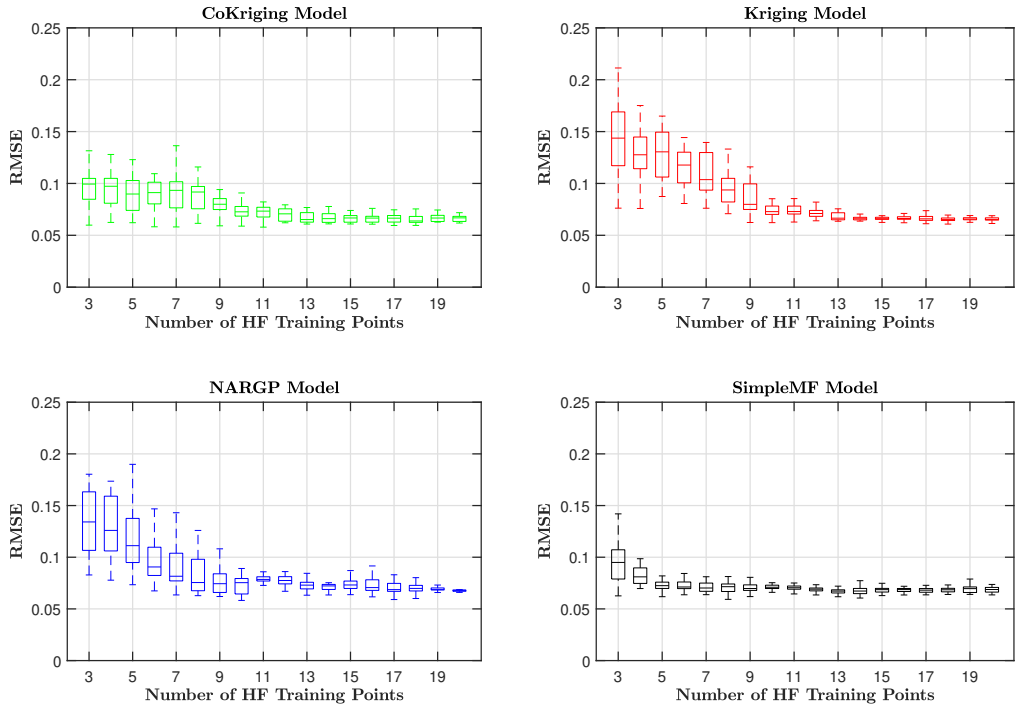

Fig. 3.17 RMSE distributions for alternative models as a function of training set size. SimpleMF achieves lowest error with minimal training data. Outliers omitted for clarity

Figure 3.18 directly compares mean RMSE across models. SimpleMF exhibits superior convergence rate, achieving asymptotic performance with approximately five to seven training points. With sufficient training data ($n \geq 14$), the more flexible CoKriging and NARGP models eventually surpass SimpleMF, confirming that the simplified approach trades ultimate accuracy for robustness in data-scarce regimes.

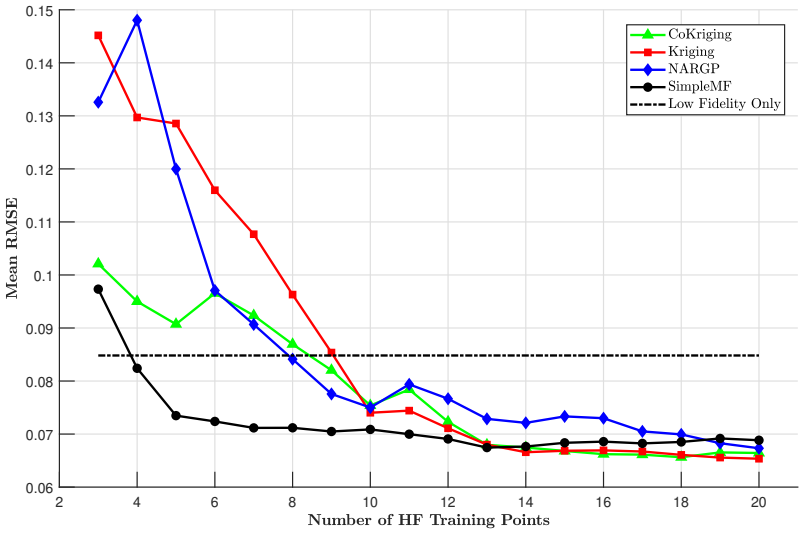

Fig. 3.18 Mean RMSE comparison demonstrating SimpleMF's rapid convergence with minimal training data

Figure 3.19 compares multi-fidelity predictions (trained with only five experimental points) against the complete experimental dataset. In contrast to the uncorrected low-fidelity results (Fig. 3.12), the SimpleMF predictions cluster tightly around the ideal correlation line, demonstrating successful capture of the underlying trends with minimal training data. The computational overhead of the multi-fidelity model itself is negligible ($O(10^{-4}$ CPU $\cdot$ h) compared to the low-fidelity simulations.

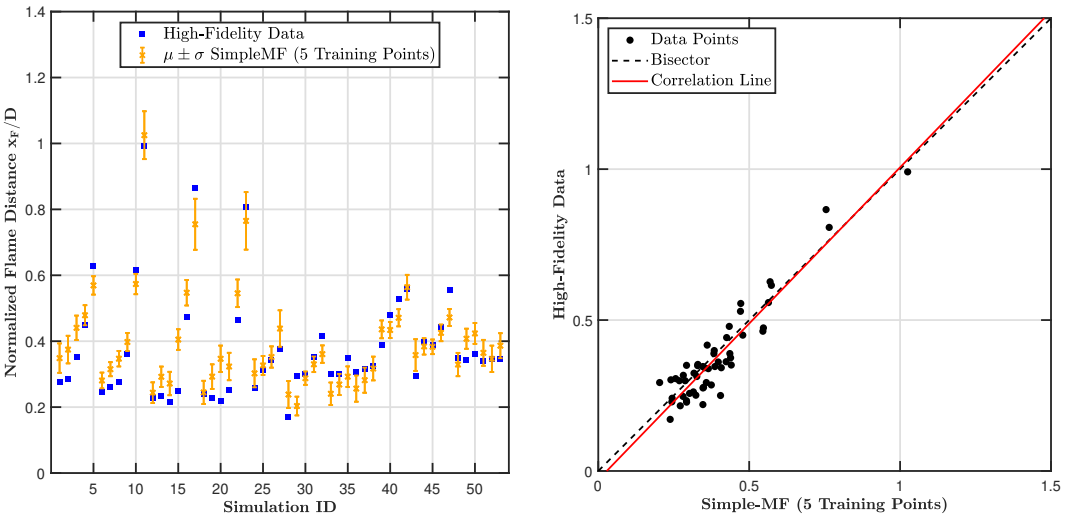

(a) Direct comparison

(b) Correlation

Fig. 3.19 Multi-fidelity predictions using five training points compared to complete experimental dataset, showing substantial improvement over uncorrected low-fidelity results. Error bars represent standard deviation across admissible training sets

3.3.8 Discussion and Transition to Chemical Reactor Networks

The SimpleMF framework successfully addresses the challenge posed at the beginning of this section: extracting maximum value from limited experimental data while maintaining computational tractability. By combining 2D RANS simulations, themselves a simplification from the 3D treatment employed for the micromix combustor, with a simple linear correlation and RBF correction, the approach achieves predictive accuracy approaching experimental measurements using training sets as small as three to five points.

The method's strengths lie in its simplicity and robustness. The minimal hyperparameter set (two linear coefficients plus RBF interpolation parameters) can be reliably estimated even when high-fidelity data are severely limited, a scenario frequently encountered in experimental combustion research where operating envelopes are constrained by safety considerations and measurement accessibility. The approach is general: while demonstrated here for flame position prediction, it applies to any combustor performance metric amenable to correlation between low- and high-fidelity models.

However, the computational cost of the 2D RANS simulations remains non-negligible at $O(7\ \mathrm{CPU}\cdot\mathrm{h})$ per operating point. This limitation becomes acute when investigating pollutant formation as nitrogen oxides. Quantitative NO_x prediction requires either substantially finer meshes to resolve reaction zone structure or Large Eddy Simulation to capture turbulence-chemistry interactions accurately [63]. Such refinements would introduce a computational burden far exceeding that which motivated the 2D simplification. In fact, as will be shown in the next section, even with the 2D simplification, the computational cost is significant when dealing with complete chemical mechanisms that include nitrogen chemistry.

This observation motivates the Chemical Reactor Network (CRN) approach developed in Section 4. By representing the combustor as a network of idealized reactors connected by mass flow controllers calibrated to match high-fidelity flow field information, CRN models achieve computational costs of $O(10^{-2}\ \mathrm{CPU}\cdot\mathrm{h})$ while retaining detailed chemical kinetics. This reduction enables rapid exploration of design parameters with NO_x prediction, complementing the flow field and flashback assessments presented in this section.

Chapter 4

Chemical Reactor Networks

Chemical Reactor Networks (CRN) constitute a modelling concept originally introduced by Bragg in the 1950s [20] to discretize combustion chambers into interconnected ideal reactors. While initially developed to evaluate fuel-air reaction rates and heat release, this chapter employs the methodology as a computationally efficient surrogate model for predicting pollutant emissions.

The approach relies on decomposing the flow field into characteristic zones, such as inner recirculation zones or flame fronts, and imposing physical constraints within specific reactors, such as the stoichiometric condition in a diffusion flame. The remaining undefined mass flow rates between reactors are treated as dimensionless parameters, which are calibrated through an optimization process that finds adequate residence times in the reactors by matching high-fidelity simulations or experimental data.

This methodology is applied to two distinct configurations: a non-premixed test case, analysed in detail in Ref. [39] and presented in Sec. 4.2, and a premixed test case, analysed in detail in Ref. [90], Ref. [89], and Ref. [42] and presented in Sec. 4.3. The general theoretical and numerical framework common to both applications is first established in Sec. 4.1.

4.1 General Numerical Framework

This section presents the theoretical framework underlying the CRN methodology employed in both case studies, encompassing the governing conservation equations, the numerical solver implementation, and the optimization strategy for parameter calibration.

4.1.1 Governing Equations

The three-dimensional flow field within a combustor can be approximated by discretizing it into a network of idealized reactors connected by mass and energy flows, where each reactor represents a specific region within the combustion chamber. For the investigations presented in this chapter, ideal gas constant volume reactor models were employed under the following assumptions: reactor walls are treated as adiabatic and chemically inert, and reactor volumes remain constant during each simulation run.

With these assumptions, the governing equations for mass conservation, species transport, and energy conservation for a generic reactor take the following forms:

Mass Conservation

$$\frac{dm}{dt} = \sum_{\text{in}} \dot{m}_{\text{in}} - \sum_{\text{out}} \dot{m}_{\text{out}} \tag{4.1}$$

where m is the mass within the reactor, and $\dot{m}_{\text{in}}$ and $\dot{m}_{\text{out}}$ represent mass flow rates through inlets and outlets, respectively.

Species Transport

$$m\frac{dY_j}{dt} = \sum_{\text{in}} \dot{m}_{\text{in}}(Y_{j,\text{in}} - Y_j) + \dot{m}_{j,\text{gen}} \tag{4.2}$$

where Y_j is the mass fraction of species j, $Y_{j,\text{in}}$ is the mass fraction of species j in the inlet stream, and $\dot{m}_{j,\text{gen}}$ is the generation rate of species j due to chemical reactions:

$$\dot{m}_{j,\text{gen}} = V\mathscr{R}_j \tag{4.3}$$

with V denoting the reactor volume and $\mathscr{R}_j$ the volumetric rate of creation of species j.

Energy Conservation

$$mc_v \frac{dT}{dt} = \sum_{\text{in}} \dot{m}_{\text{in}} \left(h_{\text{in}} - \sum_j e_j Y_{j,\text{in}} \right) - \frac{pV}{m} \sum_{\text{out}} \dot{m}_{\text{out}} - \sum_j \dot{m}_{j,\text{gen}} e_j \tag{4.4}$$

where c_v is the specific heat at constant volume, T is the temperature, p is the pressure, h_{in} is the specific enthalpy at the inlet, and e_j is the specific internal energy of species j.

In the CRN framework, the reaction rates $\mathscr{R}_j$ are calculated using direct finite-rate kinetics, in contrast to turbulent combustion models employed in CFD simulations. The use of a small set of reactors enables the integration of detailed chemical mechanisms while maintaining computational efficiency.

4.1.2 Solver and Chemical Kinetics

The reactor network simulations were performed using the open-source Cantera software package [50], which solves the governing conservation equations for each reactor in the network and performs thermochemical calculations. The governing equations were integrated in time using Cantera's built-in solver until steady-state conditions were achieved, as indicated by the stabilization of temperature, species concentrations, and mass flow rates across all reactors.

Multiple nitrogen-extended hydrogen oxidation mechanisms were evaluated across the case studies to assess the sensitivity of emission predictions to chemical kinetic complexity. The mechanisms employed include:

- **Naik mechanism** [106]: 21 species, 105 reactions
- **CRECK mechanism** [136, 142, 143]: 31 species, 203 reactions
- **Reduced Gotama mechanism** [52]: 25 species, 110 reactions (employed in the non-premixed case only)

All mechanisms incorporate the species and reaction pathways necessary to capture the principal NO_x formation routes in hydrogen combustion, including the thermal mechanism (Zeldovich), the N_2O intermediate pathway, and the NNH route.

4.1.3 Optimization Strategy

A critical challenge in CRN development is the determination of mass flow distributions between reactors. While some connections can be derived analytically from continuity equations, many interconnections are not uniquely defined and must be calibrated against high-fidelity data. Both case studies employ optimization algorithms to determine the network parameters that best reproduce experimental or CFD-predicted emissions.

The general optimization approach proceeds as follows:

1. **Network Decomposition**: The combustion chamber is discretised into zones based on characteristic flow structures (e.g., recirculation zones, flame regions, axial flow regions). Physical constraints, such as unity equivalence ratio in a diffusion flame, are imposed.

2. **Parameter Definition**: A set of dimensionless parameters is defined to characterize mass flow splits and, if necessary, volume fractions within the network. These parameters become the optimization variables.

3. **Objective Function**: An objective function f is formulated to quantify the discrepancy between CRN predictions and reference data (experimental measurements or CFD results). The objective function typically takes the form of a weighted error metric.

4. **Stochastic Optimization**: Optimization algorithms are employed to minimize the objective function and identify optimal parameter sets.

5. **Validation**: Once calibrated, the CRN is validated against operating conditions not included in the optimization dataset to assess predictive capability and generalization.

This methodology enables the construction of computationally efficient reactor networks capable of predicting emissions across a range of operating conditions

without requiring new CFD simulations for each parametric variation. The reactor network volumes may be determined through geometric approximations, swirl number analysis, or extraction from CFD flow field data, depending on the specific application requirements.

4.2 Case Study I: Non-Premixed Hydrogen Combustion with Water Injection

4.2.1 Motivation and Context

Non-premixed combustion of hydrogen offers distinct advantages in terms of operational safety, particularly in mitigating flashback phenomena. In premixed systems, the high flame speed of hydrogen combined with low-velocity regions near walls creates conditions conducive to upstream flame propagation into the premixing zone, potentially leading to catastrophic equipment damage. Non-premixed configurations avoid this risk by maintaining spatial separation between fuel and oxidizer streams until the point of combustion. However, this combustion mode can result in elevated NO_x production due to the locally high temperatures reached in near-stoichiometric flame regions. This increase in pollutants, however, could be counteracted through water injection. This case study investigates the application of CRN methodology to a laboratory-scale non-premixed swirl burner operating under ambient pressure, with the objective of predicting NO_x reduction through steam injection while maintaining constant outlet temperature through equivalence ratio adjustment.

4.2.2 Experimental Configuration

The experimental campaign was conducted at the Combustion and Optical Diagnostic Laboratory of Politecnico di Milano, utilizing a laboratory-scale non-premixed swirl burner operating under ambient pressure. The swirl burner consists of two coaxial tubes: the central tube delivers the fuel through a single-hole nozzle with an inner diameter of $d = 8$ mm positioned precisely at the burner throat, while swirling airflow is introduced through the surrounding annulus. The inner and outer radii of the annular air passage are $R_i = 7.5$ mm and $R_b = 18$ mm, respectively.

The swirling motion of air is generated through a combination of axial and tangential air inlets, with swirl intensity adjusted by varying the relative contributions of these flows. The flame is confined within a cylindrical quartz chamber with an internal diameter of 192 mm and a length of 300 mm, which provides optical access for diagnostic measurements. Exhaust gases exit through a conical hood with a 4:1 area contraction, positioned above the quartz cylinder. The experimental setup is depicted in Fig. 4.1.

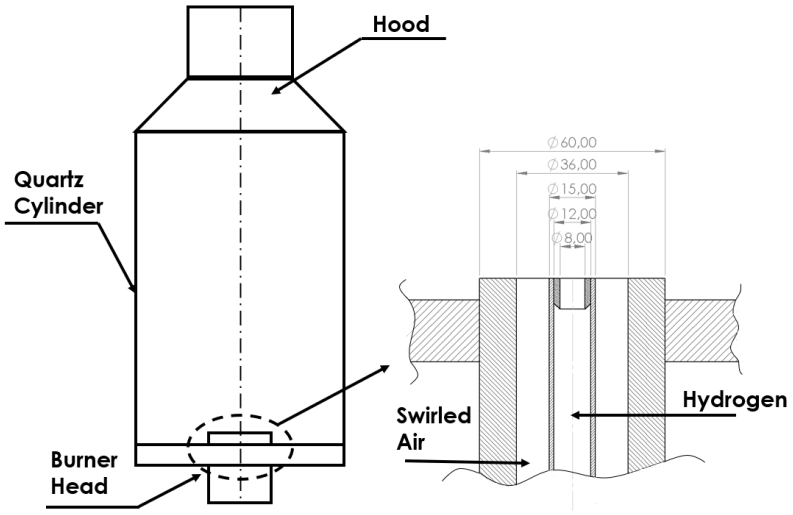

Fig. 4.1 Schematic of the burner geometry. Dimensions are in mm.

Measurements were performed at a swirl number $S = 1.0$ and an air Reynolds number $\mathrm{Re}_{\mathrm{air}} = 20{,}150$. The swirl number was calculated using the simplified definition provided by Gupta and Lilley [55], which neglects static pressure and velocity fluctuations. This computation was based on the mean axial ($\bar{v}_x$) and tangential ($\bar{v}_\theta$) velocity profiles measured via Laser Doppler Velocimetry (LDV) 1 mm downstream of the burner exit under isothermal conditions:

$$S = \frac{\int_0^{R_b} r^2 \bar{v}_x \bar{v}_\theta \, dr}{R_b \int_0^{R_b} r \bar{v}_x^2 \, dr} \tag{4.5}$$

where R_b represents the outer radius of the annular air passage and r is the radial coordinate.

The hydrogen mass flow rate was set to approximately 0.048 g/s (equivalent to 32 Nl/min), corresponding to an input thermal power of 5.8 kW, with a global

equivalence ratio of $\phi = 0.17$. The fuel jet Reynolds number, calculated at 300 K and based on the 8 mm nozzle diameter, was approximately $\mathrm{Re}_{\mathrm{fuel}} \simeq 770$. Baseline $\mathrm{NO_x}$ emissions without water injection were measured at 70 ppm (dry basis, referenced at 3% $\mathrm{O_2}$).

Nitrogen oxides (NO and $\mathrm{NO_2}$) were measured using a chemiluminescence-based technique, which relies on the gas-phase reaction between $\mathrm{O_3}$ and NO. The sampled gas was first passed through an $\mathrm{NO_2}$-to-NO converter, which reduces $\mathrm{NO_2}$ to NO, allowing the total $\mathrm{NO_x}$ concentration to be determined. Oxygen concentration was measured using a paramagnetic-based technique. All pollutant measurements are reported on a dry basis. The gas analyzer was calibrated using certified reference gases, and the standard deviation of the measured $\mathrm{NO_x}$ concentrations was approximately 4% of the mean value.

4.2.3 Computational Fluid Dynamics Framework

For this case study, two-dimensional axisymmetric RANS simulations were performed in Ansys® Fluent 2024 R1. This formulation neglects derivatives along the azimuthal direction while accounting for a tangential velocity component v_θ that varies along the axial and radial directions. The governing equations for the CFD simulations were presented in detail in Section 3.1.2.

Initially, CFD simulations were performed to validate the numerical model against experimental baseline data without water injection. Subsequently, these simulations were used to assess emissions under water-injection conditions, providing a benchmark to validate the Chemical Reactor Network (CRN) methodology through a comparative analysis of the two approaches.

Computational Domain and Boundary Conditions

The CFD domain, illustrated in Figure 4.2, excludes the air and hydrogen supply pipes. A user-defined function (UDF) was employed at the air inlet to model the swirl velocity profile as a function of the radial coordinate r and the outer radius of the annular air passage R_b, derived from the swirl number definition:

$$\frac{v_\theta}{v_x} = \frac{R_b}{r} \tag{4.6}$$

This approach ensures a swirl number of 1 at the air inlet, consistent with the experimental measurements. All walls were treated as adiabatic.

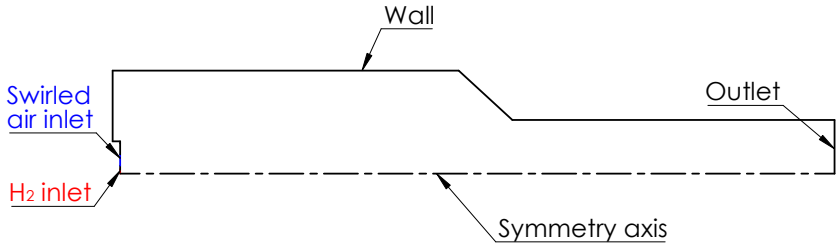

Fig. 4.2 CFD computational domain for the non-premixed swirl burner. The air and hydrogen pipes are not included in the simulation domain. All walls are treated as adiabatic. A user-defined function prescribes the swirl velocity profile at the air inlet to achieve $S = 1.0$.

Turbulence closure was achieved using the κ-ω Shear Stress Transport (SST) model [102], transporting the turbulent kinetic energy κ and the specific dissipation rate ω. Chemical reaction rates were modelled using the Eddy-Dissipation Concept (EDC) [91]. The simulations were carried out using the pressure-based solver with the coupled algorithm for pressure-velocity coupling. The pressure values at cell faces were interpolated using the PRESTO! scheme [113], while other quantities were discretised with second-order Upwind or third-order MUSCL schemes [153]. Gradient evaluation was performed using the least squares cell-based method. The complete spatial discretization scheme for all transported quantities is summarized in Table 4.1.

Table 4.1 Spatial discretization schemes employed in the CFD simulations for the non-premixed case study.

Quantity	Spatial Discretization
Pressure	PRESTO!
Density	Third-Order MUSCL
Momentum	Third-Order MUSCL
Swirl Velocity	Third-Order MUSCL
Turbulent Kinetic Energy	Second-Order Upwind
Specific Dissipation Rate	Second-Order Upwind
Species	Third-Order MUSCL
Energy	Third-Order MUSCL

Mesh Convergence Study

A systematic mesh convergence study was conducted to ensure that the numerical results were independent of grid resolution. Three progressively refined meshes were evaluated:

- **Coarse mesh**: 350 000 cells
- **Medium mesh**: 1 400 000 cells
- **Fine mesh**: 5 600 000 cells

Grid convergence was assessed by examining the variation in temperature, axial velocity, and radial velocity at an axial distance of 2 cm from the injector, a location that intersects the flame front and represents a particularly sensitive region for numerical accuracy. The results are presented in Figures 4.3, 4.4, and 4.5.

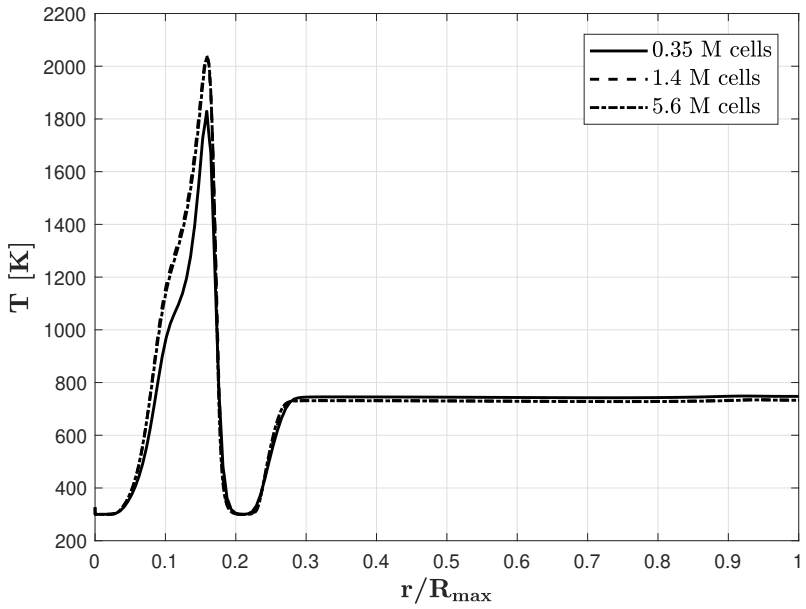

Fig. 4.3 Temperature variation across the normalized radius at an axial distance of 2 cm from the injector for three grid resolutions. The medium mesh (1.4M cells) and fine mesh (5.6M cells) show excellent agreement, demonstrating grid convergence. Naik mechanism, $Y_{H_2O} = 0\%$.

The plots consistently demonstrate that the results obtained with the 1.4 million cell mesh and the 5.6 million cell mesh are nearly indistinguishable for all three

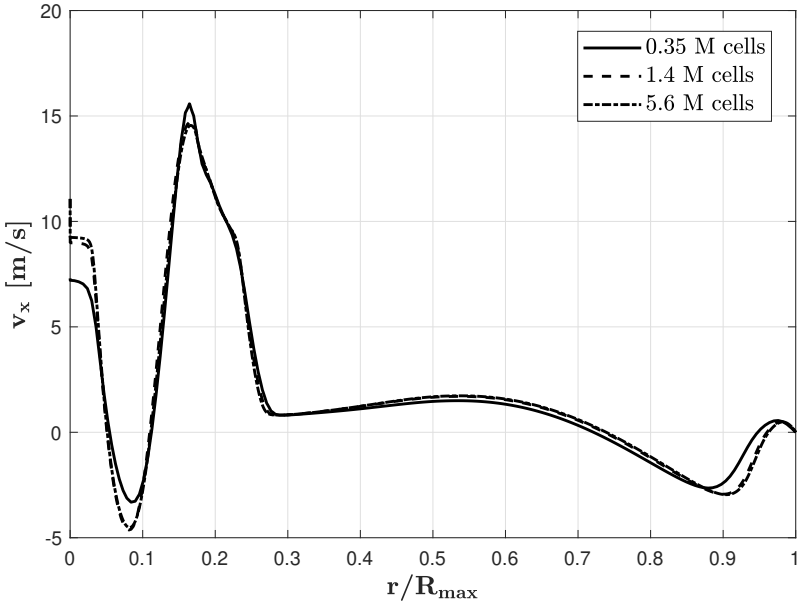

Fig. 4.4 Axial velocity variation across the normalized radius at an axial distance of 2 cm from the injector for three grid resolutions. Grid convergence is achieved with the medium mesh. Naik mechanism, $Y_{H_2O} = 0\%$.

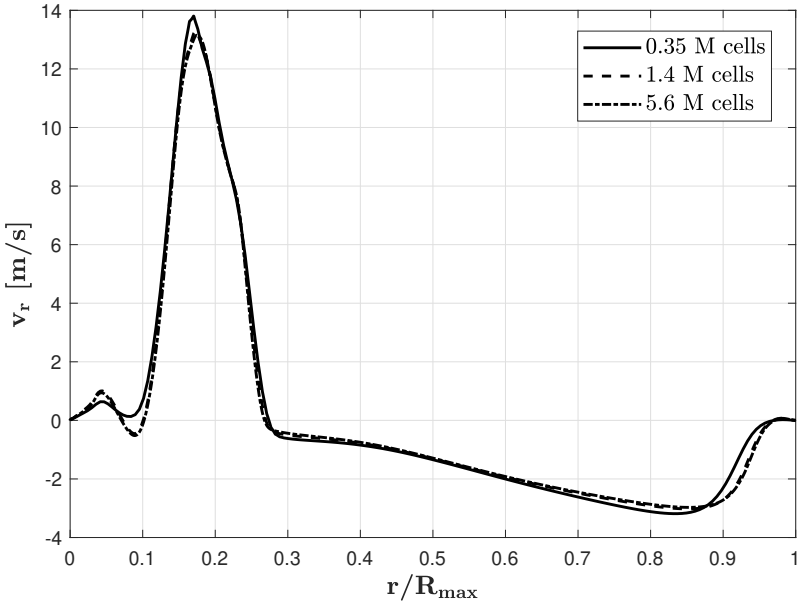

Fig. 4.5 Radial velocity variation across the normalized radius at an axial distance of 2 cm from the injector for three grid resolutions. The near-perfect overlap of the medium and fine mesh results confirms grid independence. Naik mechanism, $Y_{H_2O} = 0\%$.

variables, indicating that grid convergence has been achieved. Based on this analysis, the medium mesh was selected for all subsequent simulations, providing an optimal balance between accuracy and computational cost. Additionally, NO_x emissions were compared across meshes: the coarse mesh predicted 65.80 ppm, the medium mesh 73.86 ppm, and the fine mesh 74.91 ppm using the Naik mechanism, confirming that emission predictions are also grid-independent with the medium mesh.

Parametric Variations and Operating Conditions

To investigate the effect of steam injection on NO_x emissions while maintaining constant outlet temperature, a control strategy was implemented wherein the equivalence ratio was adjusted to compensate for the thermal effects of water addition. Specifically, the equivalence ratio was iteratively increased until the adiabatic flame temperature for the case with steam injection matched the adiabatic flame temperature for the reference case without steam addition. This approach, consistent with the methodology proposed by Concetti et al. [30], ensures that observed changes in NO_x emissions can be attributed to chemical and dilution effects rather than purely thermal effects.

Steam loadings of 5% and 10% by mass were selected for injection into the combustion chamber alongside the airflow. The equivalence ratios required to maintain constant outlet temperature are summarized in Table 4.2.

Table 4.2 Equivalence ratios required for each water mass fraction to maintain constant outlet temperature. The equivalence ratio is based on the Bilger mixture fraction formulation, including the water mass fraction in the oxidizer composition.

Y_{H_2O}	ϕ
0%	0.170
5%	0.187
10%	0.206

Comparison of Chemical Mechanisms

Three nitrogen-extended hydrogen oxidation mechanisms were evaluated in the CFD framework: the Naik mechanism (21 species, 105 reactions), the CRECK mechanism (31 species, 203 reactions), and the Reduced Gotama mechanism (25

species, 110 reactions). Temperature and velocity profiles at 2 cm from the injector were compared across mechanisms to assess their influence on the flow field structure.

Figure 4.6 shows the temperature variation across the radius for the three mechanisms alongside an isothermal flow simulation. The Naik and CRECK mechanisms exhibit nearly identical temperature profiles, with peak temperatures occurring at the stoichiometric mixture location. The Reduced Gotama mechanism shows slight deviations, predicting marginally higher temperatures in the flame core region. The isothermal simulation confirms that the velocity field modifications observed in reactive simulations are indeed consequences of combustion-induced density changes.

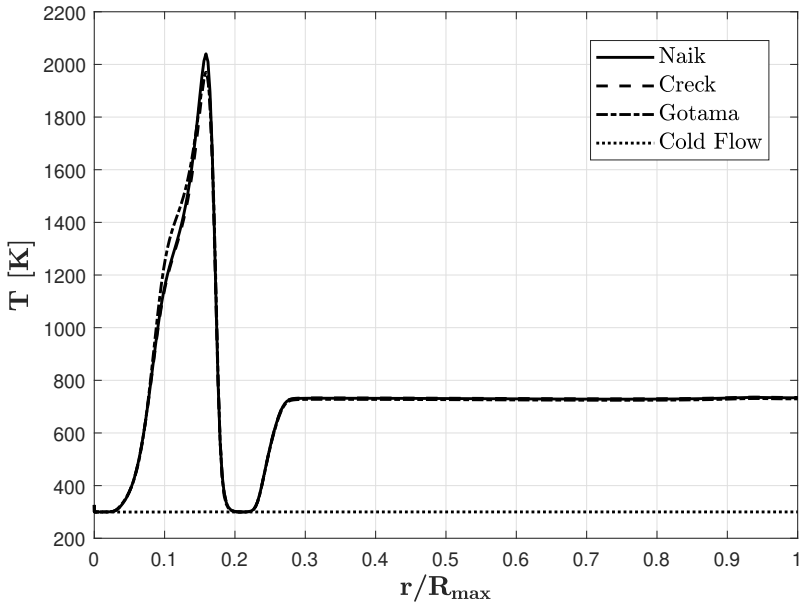

Fig. 4.6 Temperature variation across the normalized radius at an axial distance of 2 cm from the injector for different chemical mechanisms and isothermal flow. The Naik and CRECK mechanisms show excellent agreement, while the Reduced Gotama mechanism exhibits small but noticeable differences.

The axial velocity profiles, presented in Figure 4.7, reveal the jet-like structure of the hydrogen stream along the centreline, with maximum velocities exceeding 20 m/s. The reactive cases show enhanced axial velocities compared to the isothermal flow due to thermal expansion and reduced density in the high-temperature regions. The differences between chemical mechanisms are minimal, confirming that the choice of kinetic scheme has negligible impact on the global flow structure.

The radial velocity profiles, shown in Figure 4.8, illustrate the divergence of the flow as it expands from the central jet into the larger combustion chamber. The

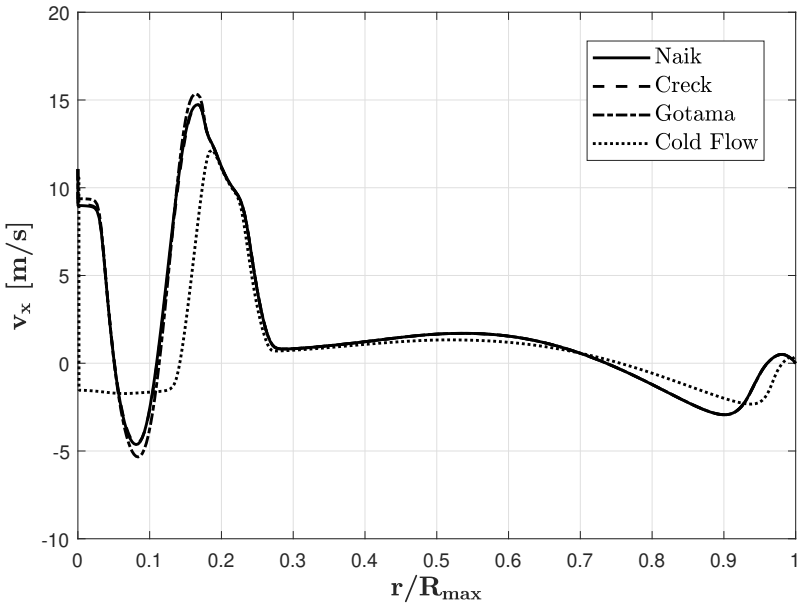

Fig. 4.7 Axial velocity variation across the normalized radius at an axial distance of 2 cm from the injector. The hydrogen jet exhibits peak velocities along the centreline. Reactive cases show enhanced velocities compared to isothermal flow due to thermal expansion. Chemical mechanism choice has minimal influence on velocity predictions.

reactive simulations exhibit stronger radial flow in the outer regions compared to the isothermal case, again attributable to density-driven expansion. The agreement between mechanisms remains excellent, validating the robustness of the flow field prediction.

Effect of Water Injection on Flow Field

A critical question for the CRN methodology is whether the flow field structure remains sufficiently stable across parametric variations to justify using a single network topology. To address this, temperature and velocity profiles were examined for different water loadings (0%, 5%, 10%) using the CRECK mechanism, with equivalence ratios adjusted according to Table 4.2 to maintain constant outlet temperature.

Figure 4.9 demonstrates that water injection successfully reduces both peak and overall flame temperatures despite the increased equivalence ratio. At $Y_{H_2O} = 0\%$, the peak temperature approaches 2400 K, while at $Y_{H_2O} = 10\%$, the peak temperature is reduced to approximately 2100 K. This 300 K reduction occurs even though the equivalence ratio increases from 0.170 to 0.206, confirming the effectiveness of

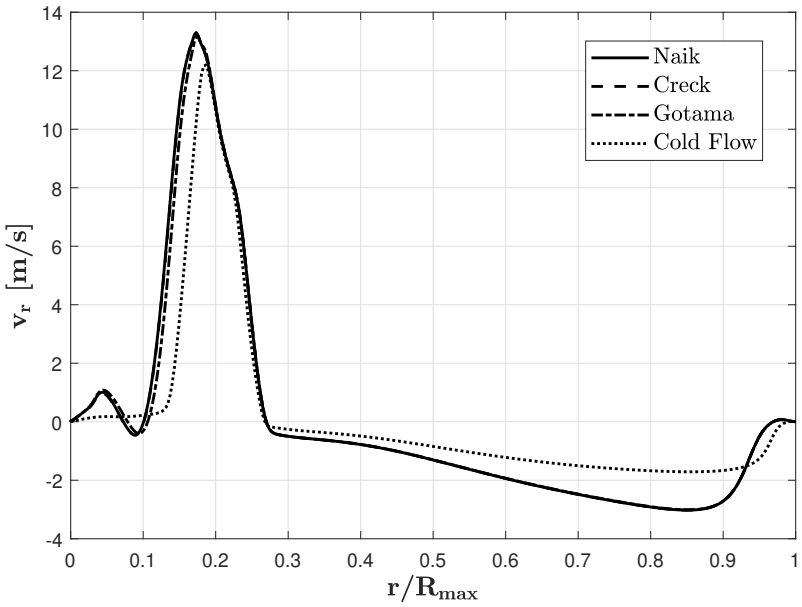

Fig. 4.8 Radial velocity variation across the normalized radius at an axial distance of 2 cm from the injector. The flow exhibits radial expansion as it transitions from the central jet to the larger chamber. Reactive cases show enhanced radial velocities due to thermal expansion. Excellent agreement is observed between all three chemical mechanisms.

water's cooling and dilution effects. The outlet temperature at approximately 850 K remains constant across all cases, validating the control strategy.

Crucially, the axial velocity profiles shown in Figure 4.10 reveal that the flow structure remains largely unaffected by water injection. The centreline jet velocities and the radial distribution of axial momentum show minimal variation across the three water loadings. This insensitivity confirms that the macroscopic flow topology, characterised by the central hydrogen jet, the surrounding swirling air annulus, and the resulting recirculation zones, persists across the operating conditions investigated.

Similarly, the radial velocity profiles in Figure 4.11 demonstrate consistent flow patterns across water loadings. The radial expansion characteristics and the location of flow reversal (indicating recirculation zone boundaries) remain essentially unchanged. This stability is critical for the CRN approach, as it implies that the reactor volume allocations and connectivity structure calibrated for the baseline case can be applied to the water injection scenarios without modification.

Temperature contour plots for the three water loadings are presented in Figures 4.12, 4.13, and 4.14, providing a visual representation of the flame structure evolution with water injection. The baseline case ($Y_{H_2O} = 0\%$) exhibits a compact,

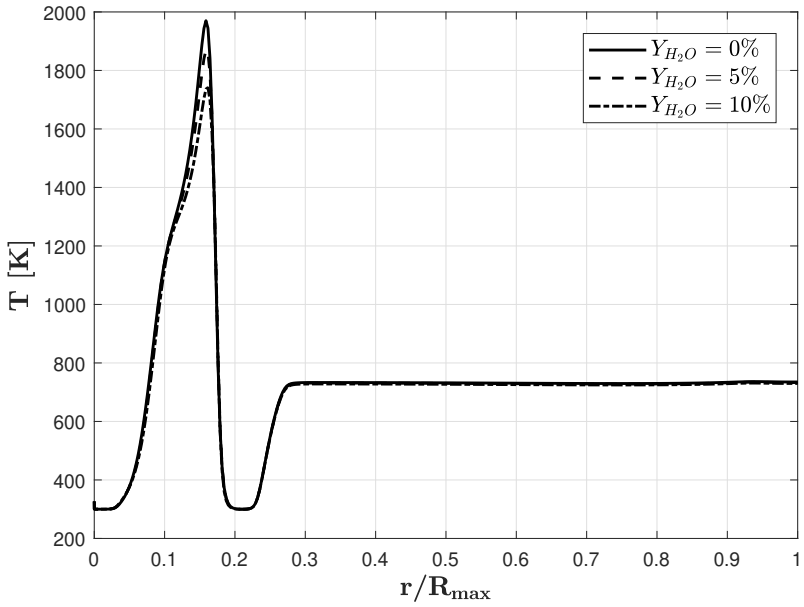

Fig. 4.9 Temperature variation across the normalized radius at an axial distance of 2 cm from the injector for different water mass fractions using the CRECK mechanism. Water injection reduces peak flame temperatures by approximately 300 K despite increased equivalence ratio, demonstrating the effectiveness of cooling and dilution effects. The constant outlet temperature (approximately 850 K) validates the control strategy.

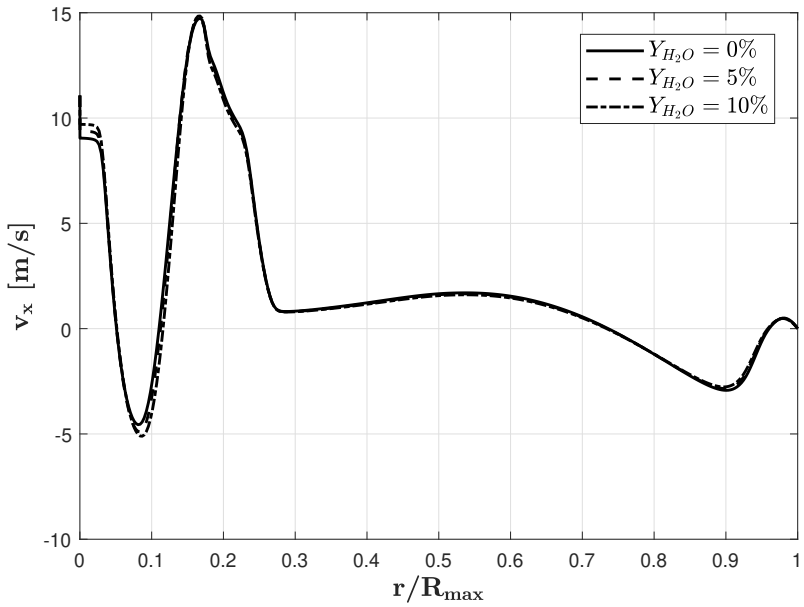

Fig. 4.10 Axial velocity variation across the normalized radius at an axial distance of 2 cm from the injector for different water mass fractions using the CRECK mechanism. The velocity field remains largely unaffected by water injection, with the hydrogen jet structure preserved across all operating conditions. This flow field stability justifies the use of a single CRN topology for all water loadings.

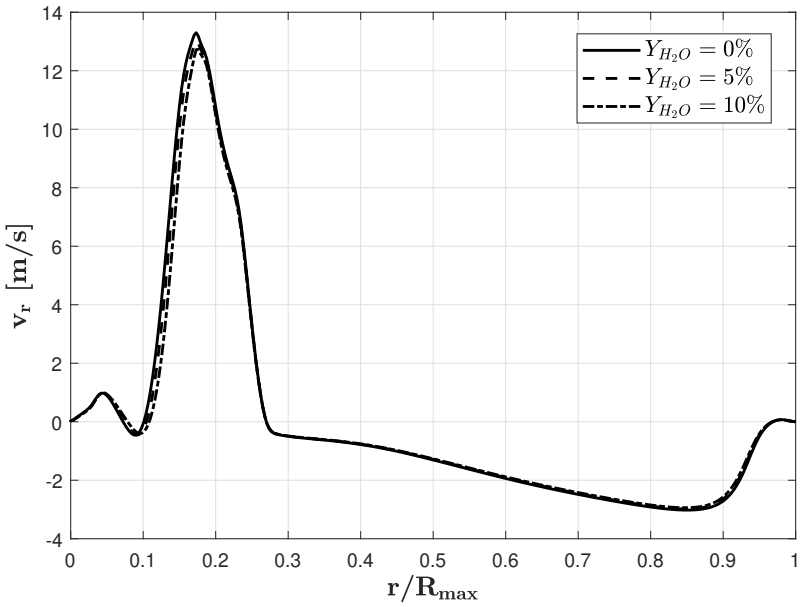

Fig. 4.11 Radial velocity variation across the normalized radius at an axial distance of 2 cm from the injector for different water mass fractions using the CRECK mechanism. The radial flow patterns and recirculation zone boundaries remain consistent across water loadings, confirming that the overall flow topology is preserved. This validates the applicability of a single CRN configuration for parametric studies.

high-temperature flame anchored near the burner exit with peak temperatures exceeding 2400 K. As water loading increases to 5% and subsequently to 10%, the high-temperature region visibly contracts and the peak temperature diminishes, while the overall flame topology remains qualitatively similar. The flame front remains anchored at approximately the same axial location, and the characteristic recirculation zones persist, confirming the suitability of a fixed CRN structure.

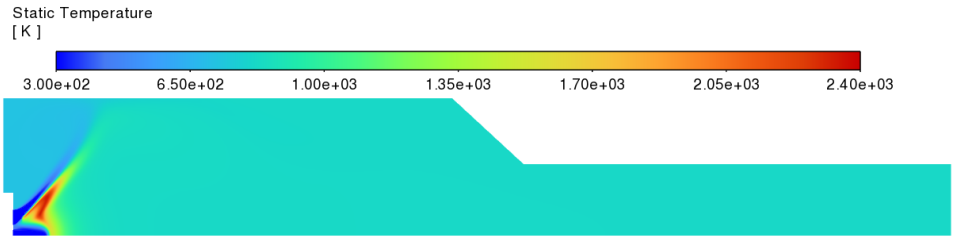

Fig. 4.12 Temperature contour plot using the Naik mechanism for the baseline case without water injection ($Y_{H_2O} = 0\%$, $\phi = 0.170$). The flame exhibits peak temperatures exceeding 2400 K in the near-stoichiometric mixing region near the burner exit.

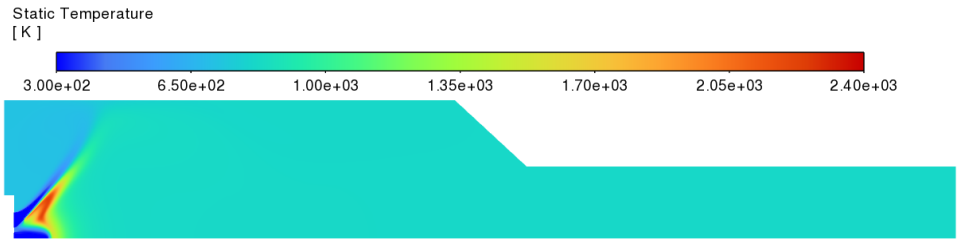

Fig. 4.13 Temperature contour plot using the Naik mechanism for 5% water injection ($Y_{H_2O} = 5\%$, $\phi = 0.187$). Water addition reduces peak temperatures despite increased equivalence ratio, demonstrating the combined cooling and dilution effects.

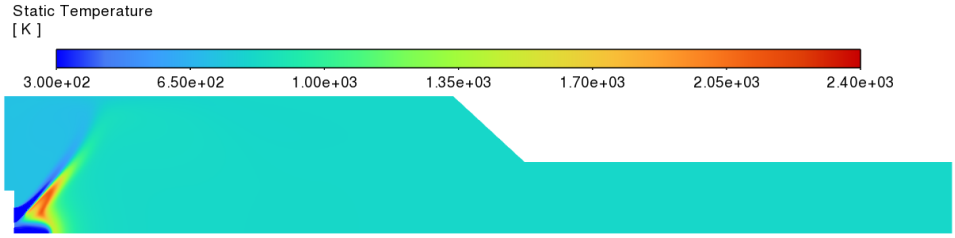

Fig. 4.14 Temperature contour plot using the Naik mechanism for 10% water injection ($Y_{H_2O} = 10\%$, $\phi = 0.206$). The high-temperature region is significantly reduced compared to the baseline, with peak temperatures approximately 300 K lower. The flame structure and anchoring location remain qualitatively similar across all water loadings.

CFD Emission Predictions

Emission predictions from the CFD simulations are reported in terms of NO_x dry ppm referenced at 3% O_2, calculated as:

$$NO_{x,\text{ppm, dry, 3\% O}_2} = (X_{NO} + X_{NO_2}) \cdot \left(\frac{X_{O_2,\text{inlet}} - X_{O_2,\text{ref}}}{X_{O_2,\text{inlet}} - X_{O_2}} \right) \cdot \frac{10^6}{1 - X_{H_2O}} \tag{4.7}$$

where X denotes mole fraction and $X_{O_2,\text{ref}} = 0.03$ represents the reference oxygen concentration.

The emission results for all three chemical mechanisms and water loadings are presented in Tables 4.3 and 4.4 for the coarse and medium meshes, respectively.

The results indicate that the CRECK mechanism provides the closest agreement with the experimental baseline value of 70 ppm at $Y_{H_2O} = 0\%$, predicting 68.00 ppm

Table 4.3 NO_x emissions predicted by CFD for different chemical mechanisms and water loadings using the 350 000 cell mesh. Values represent NO_x ppm at 3% O_2, dry basis.

CFD	Naik	Gotama	CRECK
$Y_{H_2O} = 0\%$	65.80	79.94	62.16
$Y_{H_2O} = 5\%$	21.56	28.21	18.05
$Y_{H_2O} = 10\%$	8.07	11.86	6.31

Table 4.4]

NO_x emissions predicted by CFD for different chemical mechanisms and water loadings using the 1 400 000 cell mesh. Values represent NO_x ppm at 3% O_2, dry basis.

CFD	Naik	Gotama	CRECK
$Y_{H_2O} = 0\%$	73.86	83.61	68.00
$Y_{H_2O} = 5\%$	22.79	29.10	19.64
$Y_{H_2O} = 10\%$	8.32	12.05	6.40

on the medium mesh (2.9% relative error). The Naik mechanism slightly overpredicts emissions at 73.86 ppm (5.5% relative error), while the Gotama mechanism exhibits the largest deviation at 83.61 ppm (19.4% relative error).

This discrepancy in the Gotama mechanism can be attributed to the fact that it was specifically optimized for accurately predicting laminar flame speeds of hydrogen-ammonia fuel blends rather than NO_x formation pathways. Nevertheless, despite the quantitative overprediction, the Gotama mechanism correctly captures the same qualitative trends observed with the Naik and CRECK mechanisms, particularly the significant reduction in NO_x formation with increasing water loading (from approximately 84 ppm at $Y_{H_2O} = 0\%$ to 12 ppm at $Y_{H_2O} = 10\%$), indicating that the underlying physics of water's effect on NO_x formation is represented.

Most significantly, all three mechanisms predict a dramatic reduction in NO_x emissions with water injection while maintaining constant outlet temperature. The CRECK mechanism predicts reductions of 71.1% at 5% water loading and 90.6% at 10% water loading compared to the baseline. These substantial reductions despite constant exit temperature provide definitive evidence that water's chemical and dilution effects contribute significantly to NO_x mitigation, independent of thermal pathways.

The validation of the CFD framework was carried out using emission data from the single experimental condition without water injection (70 ppm). However, the results presented on the effects of steam dilution align well with trends reported in the literature [85, 86, 88, 111], supporting the physical validity of the CFD approach for parametric studies beyond the calibration point.

The computational cost of a single CFD simulation is substantial. Estimating this cost precisely is challenging, as each refined grid run was initialized using flow fields obtained from previous simulations, and modifications such as water loading or chemical mechanism changes were applied during restarts. This strategy significantly reduced the computational expense associated with more complex chemical mechanisms and finer meshes. Nevertheless, a general estimate indicates that each simulation required approximately 10^3 CPU-hours.

4.2.4 Chemical Reactor Network Architecture

Network Topology and Physical Justification

The CRN structure for the non-premixed swirl burner comprises six Perfectly Stirred Reactors (PSRs), each representing a distinct physical region within the combustion chamber. The network topology, illustrated schematically in Figure 4.15, was designed based on analysis of the CFD temperature contours and understanding of the characteristic flow structures in swirl-stabilized non-premixed combustion.

The six reactors and their physical justifications are:

1. **PSR 1 – Hydrogen Jet Core**: This reactor represents the pure hydrogen jet emerging from the central nozzle. The high-velocity fuel stream penetrates into the combustion chamber before significant mixing with air occurs. The volume of PSR 1 is defined as the product of the inlet nozzle area and the penetration length of the hydrogen jet. The penetration length was extracted from the CFD static temperature contour (Figure 4.15), identified as the axial distance at which the hydrogen jet begins to deflect radially due to interaction with the swirling air stream. This geometric approximation ensures that PSR 1 captures the residence time of pure fuel before entering the primary reaction zone.

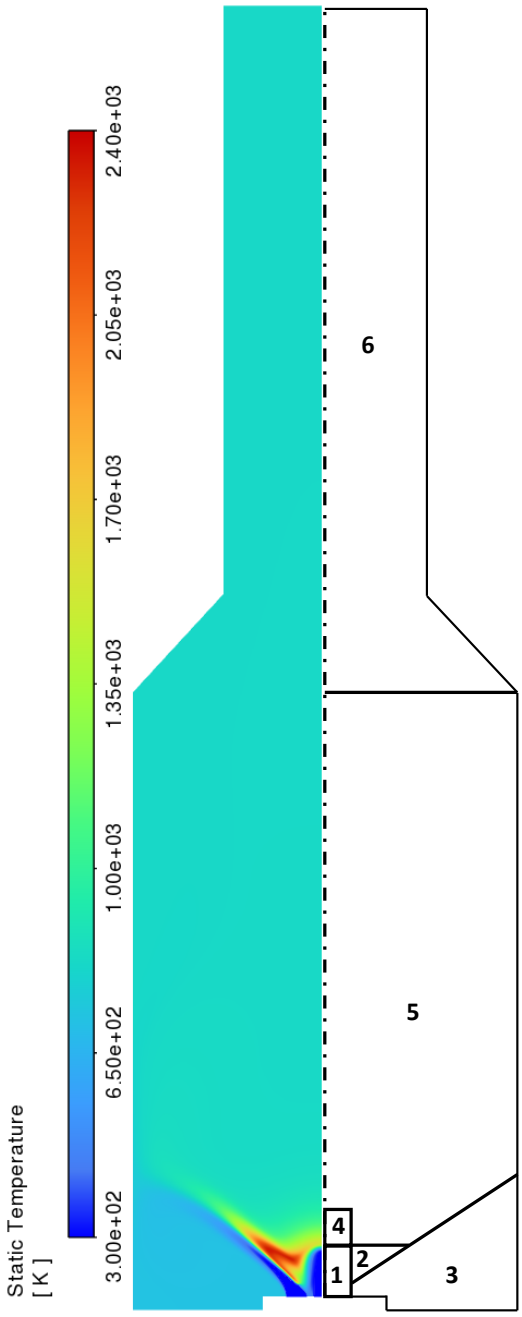

Fig. 4.15 Left panel: Static temperature contour within the combustor obtained from 2D axisymmetric CFD simulation, showing the distinct thermal zones. Right panel: Chemical reactor network model, where each numbered region (1–6) represents a Perfectly Stirred Reactor approximating the corresponding physical region. The CFD contour was used to extract the hydrogen jet penetration length and to validate the general flow structure assumed in the network decomposition.

2. **PSR 2 – Primary Flame Front**: This reactor corresponds to the high-temperature flame region where hydrogen from the jet encounters swirling air in a near-stoichiometric mixing layer. In non-premixed combustion, the flame stabilizes at the interface between fuel and oxidizer where local equivalence ratios approach unity [117], creating the peak temperatures observed in the CFD contours. PSR 2 is modelled as a truncated cone with the volume of PSR 1 subtracted. Its opening angle matches that of the outer recirculation zone (PSR 3), with an additional axial length included to account for the delayed onset of the flame downstream of the nozzle exit. This geometric construction reflects the conical expansion of the mixing layer as fuel and air interact.

3. **PSR 3 – Outer Recirculation Zone (ORZ)**: Swirl-stabilized burners generate a characteristic outer recirculation zone along the combustor walls, driven by the radial pressure gradient established by the swirling flow. This recirculation zone entrains hot combustion products and recirculates them upstream, providing thermal and chemical ignition sources that stabilize the flame. PSR 3 is approximated as a cylinder minus a cone, whose opening angle is determined by the ratio of axial to swirl velocity components as described by the swirl number $S = 1.0$. The angle α characterizing the conical exclusion is given by $\alpha = \arctan(v_x/v_\theta)$, where the velocity ratio is derived from the imposed swirl number and burner geometry. This volume allocation strategy captures the spatial extent of the recirculation region while maintaining geometric simplicity.

4. **PSR 4 – Downstream Flame Region**: As the combustion products from PSR 2 convect downstream, chemical reactions continue in a region of lower temperature and reduced chemical activity compared to the primary flame front. PSR 4 represents this post-flame zone where final oxidation of intermediate species occurs and NO_x continues to form through thermal mechanisms at elevated temperatures. The volume of PSR 4 is approximated as half that of PSR 1, reflecting the reduced spatial extent of this secondary reaction zone based on observations from the CFD temperature contours.

5. **PSR 5 – Inner Recirculation Zone (IRZ)**: In addition to the outer recirculation zone, swirl burners often exhibit a central recirculation zone on the combustor centreline downstream of the fuel jet. This IRZ is driven by the

adverse pressure gradient created by the sudden expansion from the nozzle into the larger chamber, combined with the swirl-induced vortex breakdown. The IRZ plays a crucial role in flame stabilization by recirculating hot products towards the flame root. PSR 5 is determined as the total combustor volume minus the sum of all other defined reactor volumes, ensuring that the entire combustion chamber is accounted for in the network.

6. **PSR 6 – Combustor Exit**: This reactor corresponds to the combustor volume downstream of the converging nozzle that leads to the exhaust sampling location. PSR 6 represents the final mixing and cooling region where the composition equilibrates before measurement. While chemical activity in this zone is minimal due to reduced temperatures, including this reactor ensures that the CRN predictions correspond to the same spatial location as the experimental measurements, enabling direct validation.

The complete network topology with mass flow connections is illustrated in Figure 4.16. The arrows indicate the direction and designation of mass flow rates between reactors.

Parametric Formulation and Mass Flow Allocation

While the reactor volumes are defined through geometric approximations informed by CFD, the mass flow rates connecting the reactors cannot be uniquely determined from continuity equations alone. The network contains multiple bifurcation points where flow splits must be specified. To render the system tractable for optimization, four dimensionless parameters were defined to characterize the key flow distribution decisions:

$$\mathscr{P}_1 = \frac{\dot{m}_{\mathrm{Air_I}}}{\dot{m}_{\mathrm{Air_{II}}}} \tag{4.8}$$

$$\mathscr{P}_2 = \frac{\dot{m}_B}{\dot{m}_{\mathrm{Air_I}}} \tag{4.9}$$

$$\mathscr{P}_3 = \frac{\dot{m}_F}{\dot{m}_A} \tag{4.10}$$

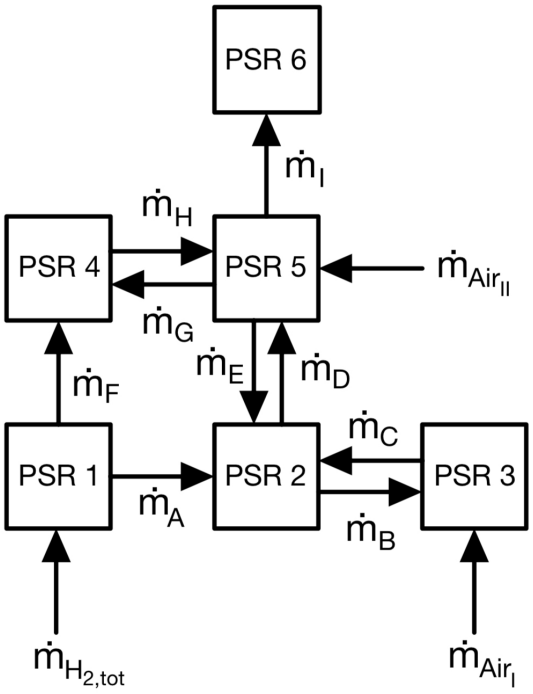

Fig. 4.16 Chemical Reactor Network topology for the non-premixed swirl burner. Arrows indicate mass flow rates between reactors, with alphabetical labels corresponding to the flow designations used in the optimization formulation. The reactor numbers correspond to the physical zones identified in Figure 4.15.

$$\mathscr{P}_4 = \frac{\dot{m}_G}{\dot{m}_I} \tag{4.11}$$

These parameters represent:

- $\mathscr{P}_1$: the split ratio of the total air mass flow rate between the stream entering PSR 3 (outer recirculation zone) and the stream entering PSR 5 (inner recirculation zone). This parameter controls the relative strength of the two recirculation zones.
- $\mathscr{P}_2$: the intensity of mass exchange between the flame, i.e. PSR 2, and the outer recirculation zone, i.e. PSR 3.
- $\mathscr{P}_3$: the ratio of mass flow exiting PSR 1 (hydrogen jet) towards PSR 4 (downstream flame) versus towards PSR 2 (primary flame). This parameter determines how much fuel participates in the primary flame front versus how much convects downstream.
- $\mathscr{P}_4$: the fraction of flow from PSR 5 (inner recirculation zone) that recirculates back to PSR 2, providing hot products for flame stabilization.

These four parameters constitute the optimization vector:

$$\mathscr{P}^{\text{opt}} = [\mathscr{P}_1, \mathscr{P}_2, \mathscr{P}_3, \mathscr{P}_4] \tag{4.12}$$

All other mass flow rates in the network can be derived analytically once these four parameters are specified, subject to the constraint of mass conservation at each reactor and a control mechanism described below.

The total hydrogen mass flow rate $\dot{m}_{\text{H}_2,\text{tot}}$ enters PSR 1, while the total air mass flow rate $\dot{m}_{\text{Air,tot}}$ is divided between two streams: $\dot{m}_{\text{Air}_\text{I}}$ enters PSR 3 and $\dot{m}_{\text{Air}_\text{II}}$ enters PSR 5. From the parameter definitions, the inlet streams are computed as:

$$\dot{m}_{\text{Air}_\text{I}} = \frac{\mathscr{P}_1}{\mathscr{P}_1 + 1} \cdot \dot{m}_{\text{Air,tot}} \tag{4.13}$$

$$\dot{m}_{\text{Air}_\text{II}} = \dot{m}_{\text{Air,tot}} - \dot{m}_{\text{Air}_\text{I}} \tag{4.14}$$

The individual flow connections (labelled $\dot{m}_A$ through $\dot{m}_I$ in Figure 4.16) are then calculated as:

$$\dot{m}_A = \frac{\dot{m}_{\mathrm{H_2,tot}}}{1+\mathscr{P}_3} \tag{4.15}$$

$$\dot{m}_B = \mathscr{P}_2 \cdot \dot{m}_{\mathrm{Air_I}} \tag{4.16}$$

$$\dot{m}_C = \frac{\dot{m}_{\mathrm{Air_I}}}{1+\mathscr{P}_2} \tag{4.17}$$

$$\dot{m}_F = \mathscr{P}_3 \cdot \dot{m}_A \tag{4.18}$$

$$\dot{m}_G = \mathscr{P}_4 \cdot \dot{m}_I \tag{4.19}$$

$$\dot{m}_I = \dot{m}_{\mathrm{Air,tot}} + \dot{m}_{\mathrm{H_2,tot}} \tag{4.20}$$

The flow rates $\dot{m}_D$ and $\dot{m}_H$ cannot be determined from the parameters alone, as they depend on $\dot{m}_E$, which is governed by a control mechanism. The mass flow rate $\dot{m}_E$ from PSR 5 to PSR 2 is determined by enforcing that PSR 2 operates as close to stoichiometric conditions as possible, reflecting the physical reality that the diffusion flame stabilizes where local mixture ratios approach $\phi = 1$. This control mechanism computes $\dot{m}_E$ to drive the equivalence ratio in PSR 2 towards unity if the inlet conditions permit, or otherwise to maximize it within the feasible range $[0,1]$.

The control algorithm employs intermediate variables based on mass fractions:

$$K_1 = \dot{m}_A \cdot Y_{\mathrm{H_2}}^{\mathrm{PSR1}} - (\dot{m}_B + \dot{m}_D) \cdot Y_{\mathrm{H_2}}^{\mathrm{PSR2}} \tag{4.21}$$

$$K_2 = \dot{m}_A \cdot Y_{\mathrm{O_2}}^{\mathrm{PSR1}} + \dot{m}_C \cdot Y_{\mathrm{O_2}}^{\mathrm{PSR3}} - (\dot{m}_B + \dot{m}_D) \cdot Y_{\mathrm{O_2}}^{\mathrm{PSR2}} \tag{4.22}$$

where $Y_{\mathrm{H_2}}^{\mathrm{PSRi}}$ and $Y_{\mathrm{O_2}}^{\mathrm{PSRi}}$ denote the mass fractions of hydrogen and oxygen in reactor i. The mass flow rate $\dot{m}_E$ is then computed as:

$$\dot{m}_E = \frac{K_1 - f_{\text{st}}^{\text{H}_2-\text{O}_2} \cdot K_2}{f_{\text{st}}^{\text{H}_2-\text{O}_2} \cdot Y_{\text{O}_2}^{\text{PSR5}} - Y_{\text{H}_2}^{\text{PSR5}}} \tag{4.23}$$

where $f_{\text{st}}^{\text{H}_2-\text{O}_2} = 0.125$ is the stoichiometric fuel-to-oxidizer mass ratio for hydrogen and oxygen. Once $\dot{m}_E$ is determined, the remaining flow rates are updated to satisfy continuity:

$$\dot{m}_D = \dot{m}_A + \dot{m}_C + \dot{m}_E - \dot{m}_B \tag{4.24}$$

$$\dot{m}_H = \dot{m}_E + \dot{m}_G + \dot{m}_I - \dot{m}_D - \dot{m}_{\text{Air}_{\text{II}}} \tag{4.25}$$

During the transient simulation as the reactors evolve towards steady state, the control mechanism is periodically invoked to recompute $\dot{m}_E$ based on the current species compositions in each reactor. This ensures that PSR 2 continuously tracks stoichiometric conditions as the chemical state evolves, and the dependent flows $\dot{m}_D$ and $\dot{m}_H$ are updated accordingly.

Optimization Procedure

The four parameters $\mathscr{P}_1, \mathscr{P}_2, \mathscr{P}_3, \mathscr{P}_4$ were calibrated using Particle Swarm Optimization (PSO) to match the experimentally measured NO_x emissions of 70 ppm at the baseline condition ($Y_{\text{H}_2\text{O}} = 0\%$, $\phi = 0.170$).

The PSO algorithm is inspired by the social behaviour of flocks of birds and schools of fish searching for food [73]. Unlike Genetic Algorithms, PSO does not require the tuning of complex genetic operators. Instead, it models candidate solutions as particles that explore the solution space, guided by their own best-known positions and the best position discovered by the entire swarm.

Each particle's motion is influenced by three components: inertia w, cognitive acceleration c_1, and social acceleration c_2. The cognitive acceleration attracts a particle towards its personal best position $\vec{p}_{\text{best},i}$, while the social component attracts it towards the global best position $\vec{g}_{\text{best}}$ found by the swarm. Particle velocities are updated at each iteration k according to:

$$\vec{v}_{k+1} = w\vec{v}_k + c_1 r_1(\vec{p}_{\text{best},i} - \vec{x}_k) + c_2 r_2(\vec{g}_{\text{best}} - \vec{x}_k) \tag{4.26}$$

where r_1 and r_2 are random variables uniformly distributed in $[0,1]$ that introduce stochastic exploration, and $\vec{x}_k$ is the position of the particle at iteration k. The balance between these terms determines the algorithm's capacity to explore new regions versus exploiting known promising areas. Particle positions are then updated as:

$$\vec{x}_{k+1} = \vec{x}_k + \vec{v}_{k+1} \tag{4.27}$$

To maintain population diversity and prevent premature convergence to local optima, boundary reflection is implemented. When a particle exceeds the prescribed search bounds, its position is reflected back into the feasible domain and the corresponding velocity component is reversed.

Adaptive inertia weighting is typically employed, wherein the inertia coefficient decreases linearly during the optimization process. This time-varying strategy facilitates global exploration during early iterations when the inertia weight is high, progressively shifting towards local exploitation as the weight decreases and the swarm converges [151].

The objective function quantifies the relative difference between CRN prediction and experimental measurement:

$$f = 1 - \left| \frac{\text{NO}_x^{\text{CRN}}}{\text{NO}_x^{\text{exp}}} \right| \tag{4.28}$$

The optimization seeks to minimize this objective by adjusting the parameter vector within prescribed bounds. Based on physical reasoning and preliminary sensitivity studies, the parameter ranges were constrained as:

- $\mathscr{P}_1 \in [0,1]$: Ensures neither recirculation zone becomes negligibly small
- $\mathscr{P}_2 \in [0,1]$: Fraction of air bypassing the primary flame
- $\mathscr{P}_3 \in [0,1]$: Fraction of fuel bypassing the primary flame
- $\mathscr{P}_4 \in [0,1]$: Recirculation fraction from IRZ

The PSO algorithm employed adaptive inertia weighting as described in Section 4.1.3, with the inertia weight w decreasing linearly from 0.9 to 0.4 over the course of the optimization. The cognitive and social acceleration coefficients were maintained at constant values $c_1 = 1.5$ and $c_2 = 1.5$ throughout. A swarm size of 10 particles was employed, with a maximum of 1000 generations permitted.

Separate optimizations were conducted for each of the three chemical mechanisms (Naik, CRECK, Reduced Gotama), as the different reaction kinetics require distinct residence time distributions to achieve the target 70 ppm emission level. To accelerate convergence, the optimization for the Naik mechanism was performed first, and the resulting optimal parameters were used as an initial guess for the CRECK and Gotama optimizations, perturbed slightly to populate the initial swarm.

The optimized parameter sets for all three mechanisms are presented in Table 4.5.

Table 4.5 Optimized CRN parameters obtained from PSO calibration for the three chemical mechanisms in the non-premixed case study. All optimizations targeted the experimental baseline of 70 ppm NO_x at $Y_{H_2O} = 0\%$.

Parameter	Naik	CRECK	Gotama	Range
$\mathscr{P}_1$	0.189	0.172	0.217	[0, 1]
$\mathscr{P}_2$	0.125	0.125	0.125	[0, 1]
$\mathscr{P}_3$	0.216	0.216	0.216	[0, 1]
$\mathscr{P}_4$	0.099	0.099	0.099	[0, 1]

Interestingly, three of the four parameters ($\mathscr{P}_2$, $\mathscr{P}_3$, $\mathscr{P}_4$) converged to identical values across all mechanisms, suggesting that these flow splits are robust features of the burner geometry and flow physics. Only $\mathscr{P}_1$, which controls the distribution of air between the outer and inner recirculation zones, varies between mechanisms. This variation reflects the different characteristic timescales for NO_x formation in each kinetic scheme: the CRECK mechanism, with its more detailed nitrogen chemistry, requires slightly less air directed to the outer recirculation zone ($\mathscr{P}_1 = 0.172$) compared to the Gotama mechanism ($\mathscr{P}_1 = 0.217$), indicating that enhanced chemical detail enables more efficient NO_x formation with shorter residence times in high-temperature regions.

The computational cost of the optimization procedure was approximately 0.5 CPU-hours on average, based on 5 restarts with random initial positions to ensure convergence to a global optimum rather than a local minimum. This cost is negligible

compared to the 10^3 CPU-hours required for a single CFD simulation, highlighting the efficiency of the CRN calibration process.

4.2.5 Results and Validation

Baseline Validation

The optimized CRN models successfully reproduced the experimental NO_x emission level of 70 ppm for all three chemical mechanisms at the baseline condition, as shown in Table 4.6. This exact agreement is expected, as the optimization explicitly targeted this value. However, the physical reasonableness of the optimized flow distributions and reactor conditions provides confidence in the network structure.

Table 4.6 CRN-predicted NO_x emissions at the baseline condition ($Y_{H_2O} = 0\%$, $\phi = 0.170$) for the three chemical mechanisms after optimization. All predictions exactly match the experimental target of 70 ppm by design.

CRN	Naik	Gotama	CRECK
$Y_{H_2O} = 0\%$	70.00	70.00	70.00

The steady-state reactor conditions for a representative case are illustrated in Figure 4.17, which depicts the temperature, pressure, equivalence ratio, and species composition in each reactor using the CRECK mechanism at $Y_{H_2O} = 0\%$. The diagram confirms that the network behaves according to design: PSR 1 contains nearly pure hydrogen at ambient temperature, PSR 2 achieves high temperatures (exceeding 2200 K) with equivalence ratio approaching unity as enforced by the control mechanism, and PSR 6 exhibits reduced temperature and the overall system equivalence ratio of 0.170. The progression of chemical species through the network is physically consistent, with hydrogen consumed primarily in PSR 2 and NO_x forming predominantly in the high-temperature reactors PSR 2 and PSR 4.

The temporal evolution towards steady state is illustrated in Figure 4.18, which shows the NO_x concentration and static temperature in PSR 6 (combustor exit) as functions of time for the CRECK mechanism at $Y_{H_2O} = 0\%$. The reactors reach steady state within approximately 0.5 seconds of simulation time, with both temperature and NO_x concentrations stabilizing to constant values. This rapid convergence confirms

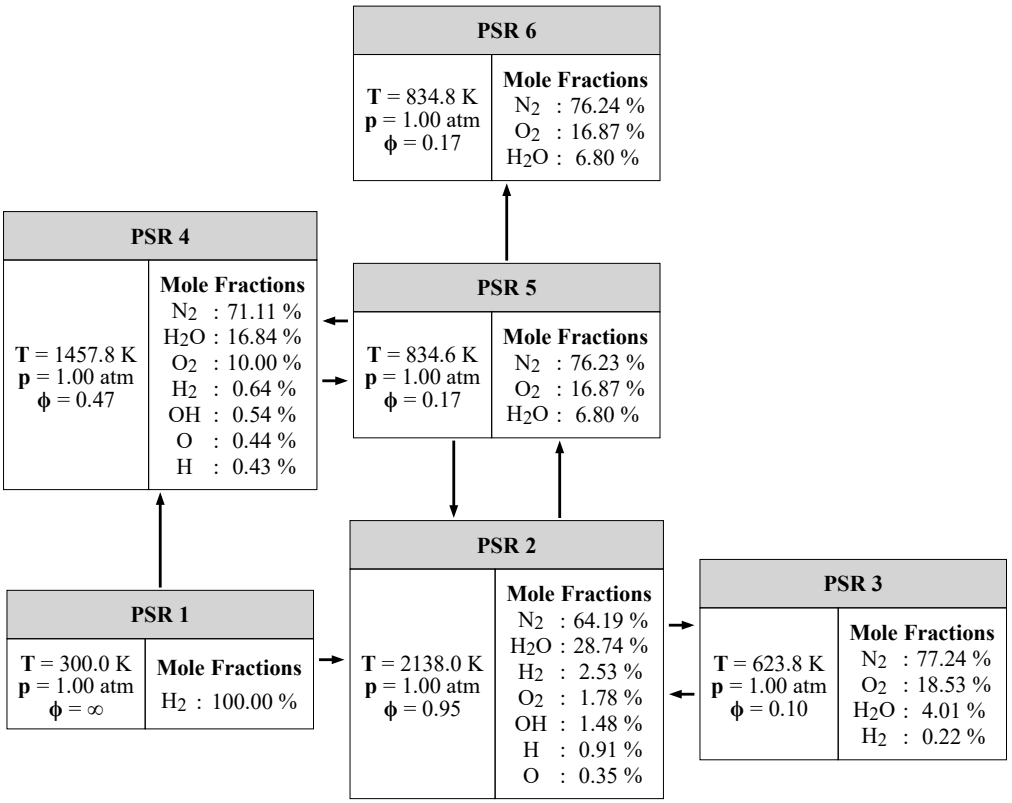

Fig. 4.17 Chemical Reactor Network at steady-state utilizing the CRECK mechanism with 0% water loading. The diagram illustrates the six interconnected reactors with their corresponding thermodynamic properties (temperature T, pressure p, equivalence ratio ϕ) and major species composition (expressed as mole fractions). Minor species with mole fractions below 0.1% are omitted for clarity. PSR 2 achieves near-stoichiometric conditions and peak temperatures exceeding 2100 K, consistent with diffusion flame physics.

numerical stability and demonstrates that the control mechanism for $\dot{m}_E$ successfully maintains PSR 2 at the target equivalence ratio throughout the transient.

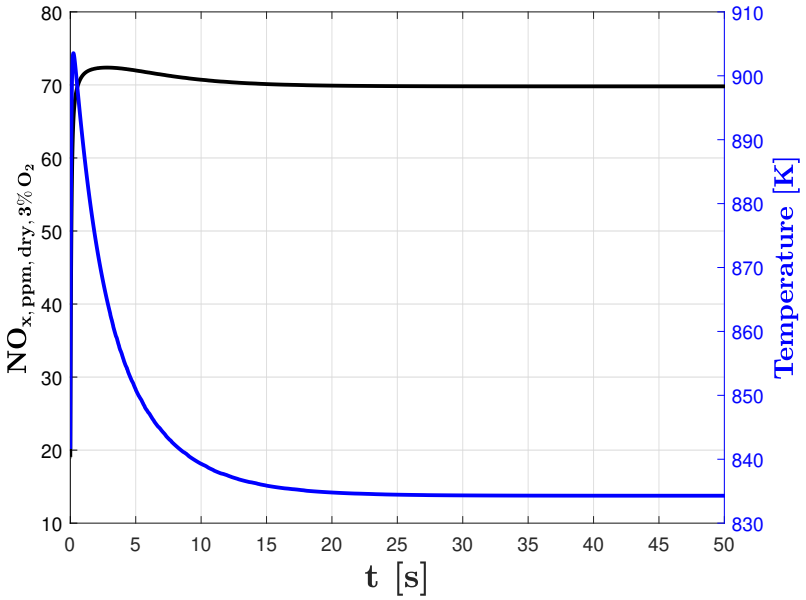

Fig. 4.18 Temporal evolution of NO_x concentration (black line) and static temperature (blue line) in PSR 6 (combustor exit) during the approach to steady state. CRECK mechanism, $Y_{H_2O} = 0\%$. The simulation reaches steady state within approximately 0.5 seconds, confirming numerical stability and appropriate residence time allocation.

Water Injection Predictions

Having calibrated the CRN against the baseline experimental measurement, the crucial test of predictive capability lies in applying the same network structure and parameters to the water injection cases. The CRN was evaluated at $Y_{H_2O} = 5\%$ and $Y_{H_2O} = 10\%$ with equivalence ratios adjusted according to Table 4.2 to maintain constant outlet temperature. Critically, no recalibration or parameter adjustment was performed; the same $\mathscr{P}^{\text{opt}}$ values from Table 4.5 were used for all water loadings.

The CRN emission predictions for all mechanisms and water loadings are presented in Table 4.7.

The CRN predictions exhibit the same qualitative trend observed in the CFD simulations: dramatic reductions in NO_x emissions with increasing water loading despite constant outlet temperature. The CRECK mechanism predicts reductions of

Table 4.7 CRN-predicted NO_x emissions for different chemical mechanisms and water loadings. Values represent NO_x ppm at 3% O_2, dry basis. The same optimized parameters from Table 4.5 were used for all water loadings without recalibration.

CRN	Naik	Gotama	CRECK
$Y_{H_2O} = 0\%$	70.00	70.00	70.00
$Y_{H_2O} = 5\%$	28.60	35.02	21.47
$Y_{H_2O} = 10\%$	11.13	18.77	7.21

69.3% at 5% water and 89.7% at 10% water compared to baseline, closely matching the CFD trends (71.1% and 90.6% reductions, respectively).

A direct comparison between CRN and CFD predictions is presented in Figure 4.19, which displays NO_x emissions for all three mechanisms and water loadings side-by-side. The CRN systematically underestimates emissions compared to CFD for the Gotama mechanism, while showing better agreement for the Naik and CRECK mechanisms. This pattern was anticipated based on the baseline calibration: since the CRN was tuned to the experimental value of 70 ppm while the Gotama CFD predicted 83.61 ppm, a systematic offset persists across all conditions.

To quantify the CRN accuracy, absolute and relative errors between CRN and CFD predictions are presented in Table 4.8.

Table 4.8 Comparison of CRN and CFD NO_x predictions showing absolute errors (CRN minus CFD, in ppm) and relative errors (percentage difference). Medium mesh CFD results (Table 4.4) are used as reference.

	Absolute Error [ppm]			Relative Error [%]		
	Naik	Gotama	CRECK	Naik	Gotama	CRECK
$Y_{H_2O} = 0\%$	-3.86	-13.61	$+2.00$	-5.2	-16.3	$+2.9$
$Y_{H_2O} = 5\%$	$+5.81$	$+5.92$	$+1.83$	$+25.5$	$+20.3$	$+9.3$
$Y_{H_2O} = 10\%$	$+2.81$	$+6.72$	$+0.81$	$+33.8$	$+55.8$	$+12.7$

At the validation condition (0% water loading), the CRECK, Naik, and Gotama mechanisms exhibit relative errors of 2.9%, 5.2%, and 16.3%, respectively, between CFD and CRN. Under water injection conditions, absolute errors evolve to $+1.83$, $+5.81$, and $+5.92$ ppm at 5% water loading, and $+0.81$, $+2.81$, and $+6.72$ ppm at 10% water loading for the CRECK, Naik, and Gotama mechanisms, respectively.

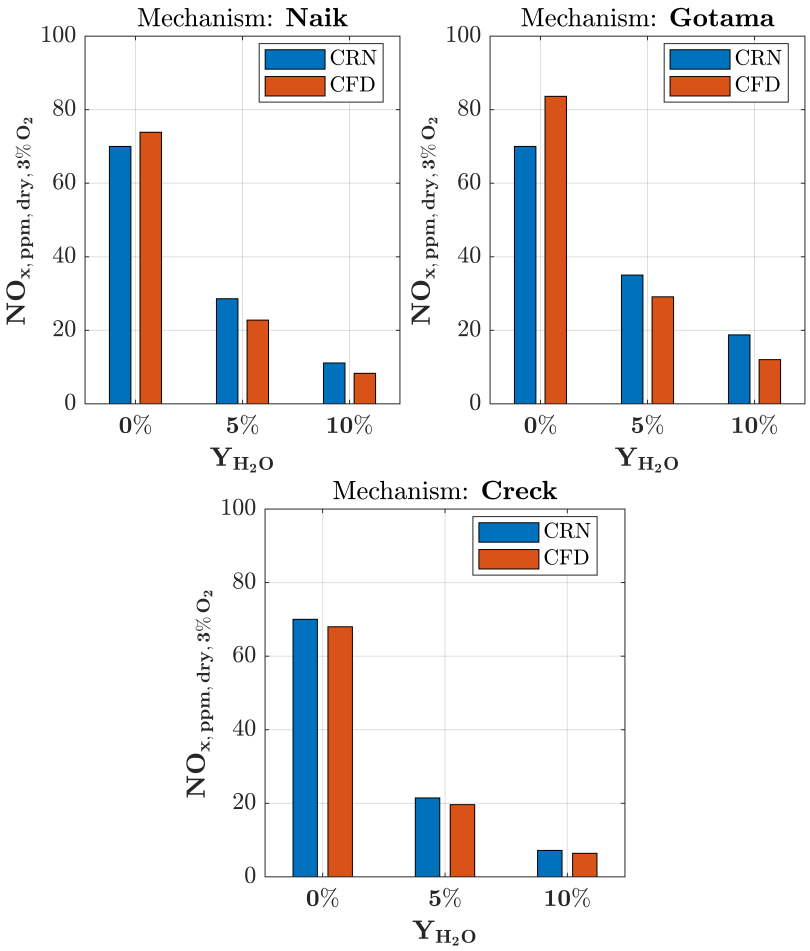

Fig. 4.19 Comparison of NO_x emissions predicted by CFD (dark orange bars) and CRN (blue bars) for the three chemical mechanisms across all water loadings. The CRN successfully captures the emission reduction trends despite being calibrated only at the baseline condition. Agreement is best for the CRECK mechanism, which also showed closest match to experimental measurements in CFD.

Importantly, the mechanism showing the best agreement at the validation condition (CRECK) maintains the most consistent performance across all water injection scenarios, with absolute errors below 2 ppm. The Naik mechanism shows intermediate behaviour, while the Gotama mechanism, which exhibited the largest deviation in the validation case, continues to show less agreement under water injection conditions. Although percentage errors generally increase with water loading due to decreasing absolute NO_x levels (making small absolute deviations appear larger in relative terms), the relative ranking of mechanism performance remains consistent.

This analysis demonstrates that the choice of chemical mechanism can be effectively guided by evaluating the CRN-CFD agreement at the validation condition, as this appears to be a reliable predictor of model consistency across different operating conditions. The CRECK mechanism, providing best agreement with both experiment and CFD at baseline, also delivers the most accurate water injection predictions.

Predictive Capability Beyond Calibration

To further demonstrate the CRN's predictive power, Figure 4.20 presents NO_x concentrations as a continuous function of water mass fraction for the CRECK mechanism. The solid black line represents the CRN predictions evaluated at closely spaced water loadings from 0% to 15%, while the discrete points show the CFD results at the three simulated conditions.

The CRN prediction curve shows excellent agreement with the discrete CFD data points across the entire range of water loadings examined. As water content increases from 0% to 15%, NO_x emissions decrease exponentially following the expected trend for thermal NO_x suppression. At 15% water loading (beyond the CFD validation range), the CRN predicts near-zero emissions, consistent with the extreme dilution and cooling effects expected at such high water fractions.

This result confirms that the reactor network serves as a computationally efficient surrogate model capable of accurately interpolating and extrapolating emission trends across water injection conditions. The ability to generate continuous emission curves at negligible computational cost enables rapid parametric exploration during the design process, a task that would be prohibitively expensive using CFD.

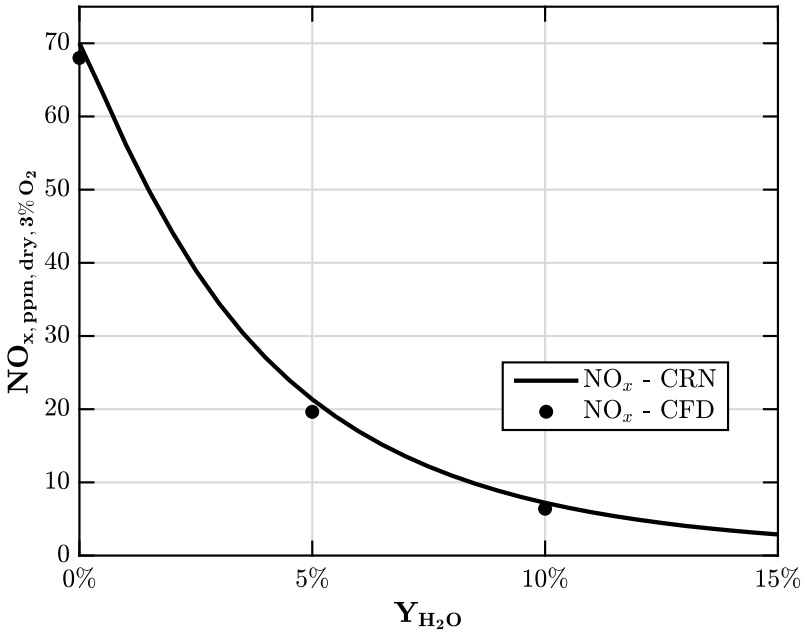

Fig. 4.20 CRN-predicted NO_x emissions (solid black line) as a continuous function of water mass fraction for the CRECK mechanism, compared with discrete CFD predictions (points). The CRN successfully captures the exponential decay trend and quantitatively matches CFD across the full water loading range. The curve demonstrates the CRN's utility as a computationally efficient surrogate model for parametric studies.

Computational Efficiency

The computational cost comparison between CFD and CRN quantifies the primary motivation for reactor network modelling. As noted earlier, each CFD simulation requires approximately 10^3 CPU-hours including initialization from previous solutions and iterative convergence to steady state. In contrast, a single CRN evaluation requires approximately 3×10^{-3} CPU-hours on a standard desktop processor (12th Gen Intel Core i7-1255U), representing a speedup of more than five orders of magnitude.

This dramatic efficiency gain enables design space exploration that would be intractable with CFD. For example, generating the continuous curve in Figure 4.20 required evaluating the CRN at 50 different water loadings, consuming less than 10 minutes of wall-clock time on a single processor core. The equivalent CFD campaign would require approximately 50,000 CPU-hours.

When the optimization cost (0.5 CPU-hours) is amortized across multiple CRN evaluations, the breakeven point occurs after just a single additional parametric variation beyond the calibration case. Given that design studies routinely explore dozens or hundreds of operating conditions, the CRN approach delivers enormous computational savings for typical applications.

Physical Insights from CRN Structure

Beyond computational efficiency, the reactor network decomposition provides physical insights into the emission formation process that are difficult to extract from CFD field data. Analysis of the steady-state reactor conditions reveals that NO_x production occurs predominantly in PSR 2 and PSR 4, the high-temperature flame regions, while the recirculation zones (PSR 3 and PSR 5) contribute minimally to NO_x formation despite their critical role in flame stabilization.

The control mechanism enforcing stoichiometric conditions in PSR 2 successfully captures the fundamental physics of diffusion flames: the reaction zone localizes where fuel and oxidizer meet in near-stoichiometric proportions, creating the peak temperatures that drive thermal NO_x formation.

The recirculation parameter $\mathscr{P}_4 = 0.099$ indicates that approximately 10% of the flow from the inner recirculation zone returns to PSR 2, providing hot products for flame anchoring. This relatively modest recirculation fraction is consistent with

the moderate swirl number ($S = 1.0$) employed in the burner design, which generates sufficient flow reversal for stabilization without creating excessive residence time in high-temperature zones.

The insensitivity of three parameters ($\mathscr{P}_2$, $\mathscr{P}_3$, $\mathscr{P}_4$) to chemical mechanism choice suggests that these flow splits are primarily governed by fluid mechanics (burner geometry, swirl intensity, Reynolds number) rather than chemistry. Only the air distribution between recirculation zones ($\mathscr{P}_1$) varies with mechanism, as different kinetic schemes require different residence time distributions to achieve the same overall emission level.

Limitations and Applicability

Several limitations of the CRN approach merit discussion. First, the network was calibrated using only a single experimental data point at the baseline condition without water injection. While subsequent validation against CFD for water injection cases demonstrates good predictive capability, true experimental validation across the water loading range would strengthen confidence in the predictions. The close agreement between CRN and CFD trends provides indirect support, given that the CFD framework itself was validated against the experimental baseline and shows physically consistent behaviour matching literature observations on water/steam dilution effects.

Second, the assumption of perfectly stirred reactors neglects spatial gradients within each zone. In reality, temperature and composition vary continuously through the combustion chamber, and reaction rates are nonlinear functions of these local conditions. The success of the model despite this approximation suggests that volume-averaged conditions provide sufficient fidelity for global emission prediction, even if local details are not resolved.

Third, the network topology and volume allocations were determined through geometric approximations informed by a single CFD solution. Changes in operating conditions that significantly alter the flow structure (e.g., dramatic changes in swirl number, Reynolds number, or burner geometry) would likely require re-evaluation of the reactor definitions. However, the demonstrated insensitivity to water injection, which modifies temperatures, densities, and chemical composition while preserving

flow topology, suggests robustness within a reasonable operating envelope around the design point.

Finally, the CRN cannot predict phenomena that depend on spatial resolution, such as local extinction and flame liftoff. For these applications, CFD or higher-fidelity methods remain necessary. The CRN is best suited for steady-state emission prediction and preliminary design optimization where computational efficiency is paramount.

Summary of Case Study I

The non-premixed swirl burner case study demonstrates that Chemical Reactor Network methodology can successfully predict NO_x emission trends across parametric variations with minimal calibration data and negligible computational cost. A reactor network, physically motivated by the characteristic flow structures in swirl-stabilized combustion, was calibrated against a single experimental measurement using four optimized flow split parameters. The calibrated network accurately reproduced CFD-predicted emission reductions of 70–90% under water injection conditions despite being trained only at the baseline, validating the approach for parametric design studies.

The computational efficiency gain of orders of magnitude compared to CFD enables rapid design space exploration, while the physically interpretable network structure provides insights into emission formation mechanisms. The methodology's success depends on two key factors: careful decomposition of the combustion chamber into zones corresponding to distinct physical regions, and stability of the flow topology across the operating conditions of interest. When these conditions are met, CRN modelling provides a powerful tool for preliminary combustor design and optimization in the context of hydrogen aviation propulsion systems.

4.3 Case Study II: Lean Premixed Swirl-Stabilized Hydrogen Combustor

4.3.1 Motivation and Context

Lean premixed combustion represents an alternative approach to NO_x reduction in hydrogen-fuelled gas turbines, offering the potential for lower emissions compared to non-premixed configurations through reduced peak flame temperatures. By thoroughly mixing fuel and oxidizer upstream of the combustion zone and operating at fuel-lean conditions, premixed systems can achieve more uniform temperature distributions and avoid the near-stoichiometric hot spots characteristic of diffusion flames. However, this approach introduces significant technical challenges, particularly the risk of flashback due to hydrogen's exceptionally high flame speed and wide flammability limits.

The AHEAD combustor investigated in this case study employs a lean premixed swirl-stabilized configuration specifically designed to mitigate flashback risk while maintaining low emissions. As described in detail in Section 3.3.3, the combustor incorporates several features: a premixing section, dilution air injection near the wall to create a protective low-equivalence-ratio boundary layer, and swirl-induced recirculation zones that anchor the flame downstream of the mixing region.

This case study demonstrates the application of CRN methodology to a premixed configuration where the fundamental combustion mechanism and flow topology differ substantially from the non-premixed burner examined in Section 4.2. A distinguishing feature of the approach presented here is a volume allocation strategy based on swirl number estimation and geometric analysis, which reduces the requirement for extensive preliminary CFD simulations while maintaining physical consistency with the underlying flow structure. This advancement addresses a critical limitation of conventional CRN methodologies that rely extensively on CFD for network calibration, thereby further enhancing computational efficiency.

4.3.2 Experimental Configuration and Operating Conditions

The geometric configuration, operating principles, and boundary conditions for the AHEAD combustor were presented comprehensively in Section 3.3.3. The present

section focuses on the experimental NO_x emission measurements that serve as the validation dataset for the CRN methodology.

NO_x emissions were measured by Reichel et al. downstream of the burner outlet and transported through a heated tube to a cold steam trap to remove humidity, with uncertainties of $\pm 0.8\,\%$ for NO and of $\pm 1.5\,\%$ for NO_2 [126]. The experimental dataset [112] encompasses a comprehensive range of operating conditions at atmospheric pressure spanning:

- Inlet air temperature: 313 K to 693 K
- Inlet air mass flow rate: 116 kg/h to 255 kg/h
- Equivalence ratio: 0.30 to 0.96

The complete experimental emission dataset is presented in Table 4.9, reporting NO_x concentrations in parts per million by volume on a dry basis, corrected to 15% O_2 concentration.

Due to the availability of a more complete dataset across a wide range of equivalence ratios, the operating condition with an inlet air mass flow rate of 130 kg/h and an inlet air temperature of 623 K was selected as the primary reference case for CRN calibration, as this condition provides eight experimental data points spanning equivalence ratios from 0.30 to 0.96.

4.3.3 Swirl Number Estimation and Volume Allocation

A critical distinction of this case study from the non-premixed configuration is the volume allocation strategy, which relies on swirl number estimation rather than direct extraction from CFD flow fields. This approach enables CRN construction without requiring high-fidelity simulations, significantly reducing the upfront computational investment.

Swirl Number Estimation

The swirl number, a dimensionless parameter quantifying the degree of rotation imparted to the axial flow, critically influences recirculation zone structure and

Table 4.9 Experimental NO_x emission measurements for the AHEAD hydrogen combustor under various operating conditions. Emissions are reported as parts per million by volume on a dry basis, corrected to 15% O_2 concentration. Data extracted from Ref. [112].

ID	$\dot{m}_{air}$ [kg/h]	T_{air} [K]	ϕ [-]	NO_x [ppm, 15% O_2, dry]
1	255	313	0.30	1.83
2	255	313	0.40	2.32
3	255	313	0.50	3.61
4	255	313	0.60	6.22
5	180	453	0.30	3.56
6	180	453	0.40	3.69
7	180	453	0.50	5.94
8	180	453	0.60	9.39
9	130	623	0.30	3.93
10	130	623	0.40	5.28
11	130	623	0.50	10.12
12	130	623	0.60	26.13
13	130	623	0.70	69.40
14	130	623	0.80	182.33
15	130	623	0.87	212.07
16	130	623	0.96	305.54
17	116	693	0.30	3.90
18	116	693	0.40	6.75
19	116	693	0.50	16.64

extent. Following the formulation of Beér and Chigier [15], the swirl number is defined as:

$$S = \frac{G_\theta}{G_x \cdot R} \tag{4.29}$$

where G_θ represents the axial flux of angular momentum, G_x denotes the axial thrust, and R is a characteristic radius. These quantities are formally expressed as:

$$G_\theta = \int_0^R (Wr)\rho U \, 2\pi r \, dr \tag{4.30a}$$

$$G_x = \int_0^R U\rho U \, 2\pi r \, dr + \int_0^R p \, 2\pi r \, dr \tag{4.30b}$$

where U, W, ρ, and p denote the axial velocity, tangential velocity, density, and static pressure, respectively, in any cross-section of the jet.

For guide-vane cascade configurations in radial flow arrangements, a simplified approach provides acceptable accuracy for preliminary design. The velocity ratio at the swirler exit can be expressed as:

$$\sigma = \frac{W_1}{V_1} \tag{4.31}$$

where W_1 and V_1 denote the tangential and radial velocities at the swirler exit. These quantities are related to the air mass flow rate through the radial swirler ($\dot{m}_{\text{swirl}}$), the mixing tube radius (R_{mix}), air density (ρ_{air}), and the axial width of the swirler channels (B):

$$W_1 = \frac{G_\theta}{\dot{m}_{\text{swirl}} R_{\text{mix}}} \tag{4.32a}$$

$$V_1 = \frac{\dot{m}_{\text{swirl}}}{\rho_{\text{air}} \, 2\pi R_{\text{mix}} B} \tag{4.32b}$$

Substituting Equations (4.32) into Equation (4.31) and rearranging yields:

$$G_\theta = \sigma \frac{\dot{m}_{\text{swirl}}^2}{\rho_{\text{air}} 2\pi B} \tag{4.33}$$

The tangential-to-radial velocity ratio $\sigma = 0.8$ was adopted based on validated CFD-PIV comparisons for this combustor configuration. As discussed in Section 3.3.5, this value was established by matching two-dimensional axisymmetric RANS predictions with experimental particle image velocimetry measurements at a representative operating condition, and subsequently applied across the full range of simulated conditions.

As a first approximation suitable for preliminary design, the flow inside the mixing conduit is assumed to be fully developed, allowing the pressure term in Equation (4.30) to be neglected. Under this assumption, the axial thrust at the mixing tube exit can be expressed as:

$$G_x = \frac{\dot{m}_{\text{mix}}^2}{\rho_{\text{mix}} A_{\text{mix}}} \tag{4.34}$$

where $\dot{m}_{\text{mix}}$ is the total mass flow rate at the mixing tube exit, ρ_{mix} is the density of the premixed fuel-air mixture, and $A_{\text{mix}} = \pi R_{\text{mix}}^2$ is the cross-sectional area of the mixing conduit.

Assuming conservation of the axial flux of angular momentum G_θ throughout the mixing tube, the swirl number at the mixing conduit exit can be estimated by combining Equations (4.29), (4.33), and (4.34):

$$S = \sigma \frac{\rho_{\text{mix}}}{\rho_{\text{air}}} \frac{\dot{m}_{\text{swirl}}^2}{\dot{m}_{\text{mix}}^2} \frac{R_{\text{mix}}}{2B} \tag{4.35}$$

This formulation enables the estimation of swirl number as a function of operating conditions, including variations in inlet temperature, equivalence ratio, and inlet air flow rate, without requiring dedicated flow field simulations. While this approach neglects the decay of G_θ along the mixing conduit and does not account for the pressure contribution near the sudden expansion, these simplifications are considered acceptable for preliminary design applications.

4.3.4 Chemical Reactor Network Architecture

Network Topology and Physical Justification

The CRN structure for the AHEAD premixed combustor comprises two inlet reservoirs representing the mixing tube and five perfectly stirred reactors discretizing the combustion chamber into distinct zones based on characteristic flow structures. The network topology is illustrated in Figure 4.21, which displays both the experimental flow field obtained from Particle Image Velocimetry and the corresponding reactor network decomposition.

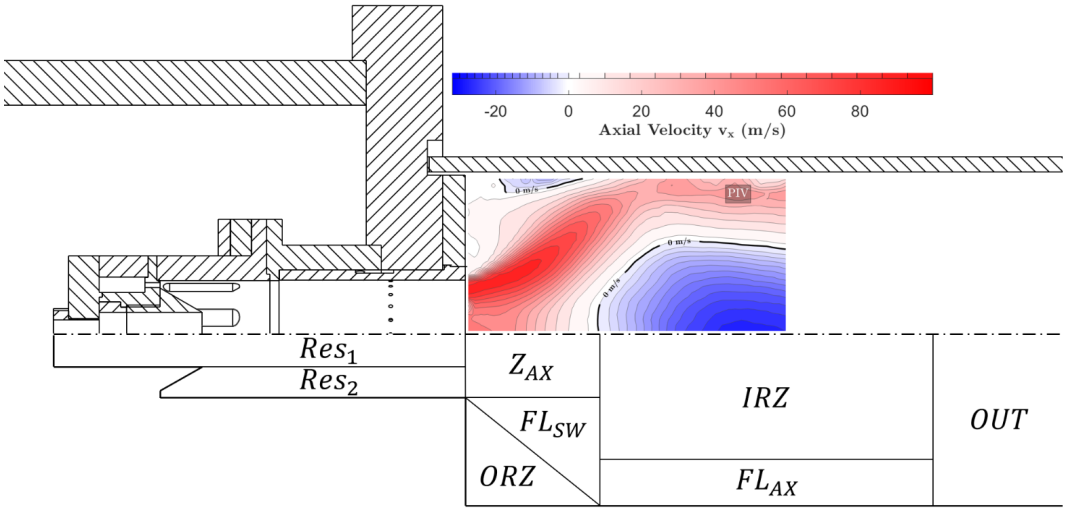

Fig. 4.21 Combustion chamber discretization employed in the CRN model. Top: experimental axial velocity field obtained from Particle Image Velocimetry (PIV) measurements, illustrating the characteristic flow structures of the swirl-stabilized combustor. The PIV data clearly shows the central inner recirculation zone (IRZ) with negative axial velocities, the outer recirculation zone (ORZ) along the walls, and the axial flow region (Z_{AX}) downstream of the mixing tube. Bottom: corresponding decomposition into the two upstream reservoirs (Res_1 and Res_2) and the five reactor zones representing the mixing tube and combustion chamber, respectively.

The mixing tube is represented by two ideal reservoirs that emulate the upstream air-fuel premixing prior to entering the main combustion chamber. This dual-reservoir approach enables representation of non-uniformities in the fuel-air distribution that may arise from the discrete hydrogen injection ports and the complex three-dimensional flow patterns within the mixing section.

Each reservoir is parameterized by its equivalence ratio (ϕ_1 and ϕ_2), defined according to the Bilger formulation based on elemental mass fractions, and air mass flow rate ($\dot{m}_{\text{air},1}$ and $\dot{m}_{\text{air},2}$). The splitting of the total air and fuel streams between these two reservoirs is governed by optimization parameters θ_1 and θ_2 described in Section 4.3.4.

The main combustion chamber is discretised into six zones based on the characteristic flow structures visible in the PIV measurements:

1. **Axial Flow Zone ($\mathbf{Z_{AX}}$)**: Region immediately downstream of the mixing tube where the premixed stream enters the combustion chamber and the flame front most frequently occurs. This zone experiences predominantly axial flow towards the combustor core.

2. **Swirling Flame Zone ($\mathbf{FL_{SW}}$)**: Annular region characterised by intense chemical reactions and significant radial velocity components induced by the upstream swirl generator. This zone surrounds the central inner recirculation zone and serves as a primary reaction region where substantial heat release and NO_x formation occur. The swirling motion enhances mixing and prolongs residence time, promoting complete combustion.

3. **Inner Recirculation Zone (IRZ)**: Central vortex region characterised by reverse axial flow, clearly visible in the PIV data as the zone with negative axial velocities along the combustor centreline. The IRZ is driven by vortex breakdown induced by the strong swirl, creating an adverse pressure gradient that reverses the flow direction. This recirculation zone entrains hot combustion products and transports them upstream, providing thermal and chemical ignition sources essential for flame stabilization. The IRZ acts as a "pilot flame" that ensures continuous ignition of the incoming fuel-air mixture.

4. **Axial Flame Zone ($\mathbf{FL_{AX}}$)**: Region downstream of the FL_{SW} zone, bounded by the combustor wall and the IRZ. The flow in this zone is predominantly axial with reduced swirl intensity compared to FL_{SW}. Chemical reactions continue in this region, particularly the oxidation of intermediate species and thermal NO_x formation at elevated temperatures, though reaction rates are lower than in the primary flame zones (Z_{AX} and FL_{SW}).

5. **Outer Recirculation Zone (ORZ)**: Peripheral recirculation zone adjacent to the chamber walls, visible in the PIV data as regions of flow reversal near the outer boundary. The ORZ is generated by the radial pressure gradients established by the swirling flow and the sudden expansion from the mixing tube into the larger combustion chamber. This zone enhances mixing between combustion products and fresh reactants and contributes to overall flame stabilization, though its role is secondary to the IRZ in this particular combustor configuration.

6. **Exit Zone (OUT)**: Downstream reactor that collects the outflow from both the IRZ and the FL_{AX} zone before discharging to the exhaust.

The complete mass flow connectivity between reservoirs and reactors is illustrated in Figure 4.22.

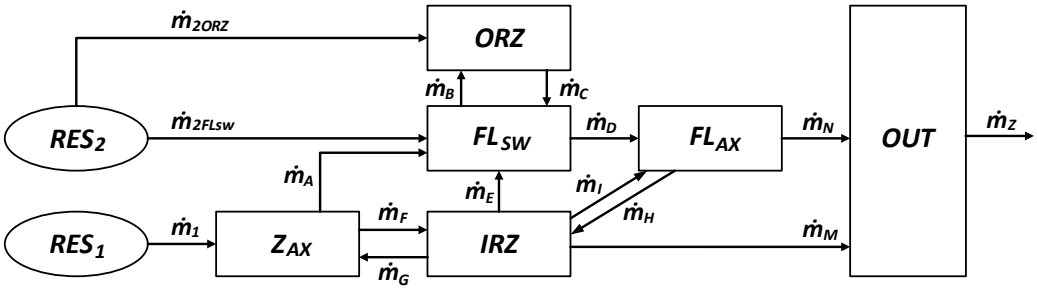

Fig. 4.22 Schematic representation of the chemical reactor network structure for the AHEAD premixed combustor. The mixing tube is modelled by two reservoirs (Res_1 and Res_2), while the combustion chamber is discretised into six interconnected perfectly stirred reactors (Z_{AX}, FL_{SW}, IRZ, FL_{AX}, ORZ, OUT). Arrows indicate mass flow connections between reactors, with alphabetical labels corresponding to the flow designations used in the mathematical formulation. The dashed lines represent recirculation pathways that provide thermal and chemical feedback for flame stabilization.

Volume Allocation Strategy

The volumes of the five reactor zones representing the combustion chamber were determined through geometric analysis and fluid dynamic approximations, integrating experimental observations, direct CAD measurements, and simplified analytical expressions for flow-dependent quantities.

The volume of the axial flow zone was calculated assuming a cylindrical geometry with cross-sectional area equal to that of the mixing tube and an axial length corresponding to the region downstream of the mixing tube where the flame front was documented by Reichel et al. [124]:

$$V_{\mathrm{Z_{AX}}} = \pi R_{\mathrm{mix}}^2 L_{\mathrm{AX}} \tag{4.36}$$

where $R_{\mathrm{mix}} = 17$ mm represents the mixing tube radius and $L_{\mathrm{AX}} = 51$ mm denotes the axial distance.

The volume of the outer recirculation zone was estimated by approximating this region as a solid of revolution about the combustor axis, with the generating profile idealized as a right-angled triangle. One leg corresponds to the radial extent of the recirculation zone (the difference between the combustion chamber radius $R_{\mathrm{cc}} = 52.5$ mm and the mixing tube radius R_{mix}). The other leg is determined by the characteristic deflection angle γ that the swirling flow creates relative to the combustor axis upon exiting the mixing tube. For the present analysis, this angle is assumed to be directly related to the swirl number through the relation [144]:

$$\tan(\gamma) = S \tag{4.37}$$

Under these assumptions, the volume of the outer recirculation zone was formulated as:

$$V_{\mathrm{ORZ}} = \pi (R_{\mathrm{cc}} - R_{\mathrm{mix}})^2 \cot(\gamma) \left(R_{\mathrm{mix}} + \frac{2}{3}(R_{\mathrm{cc}} - R_{\mathrm{mix}}) \right) \tag{4.38}$$

The volume of the swirling flame zone varies due to the interaction between the swirling flow and chamber confinement. A conditional formulation was implemented based on the relationship between the deflection angle γ and a geometric parameter α defined as:

$$\alpha = \arctan\left(\frac{R_{\mathrm{cc}} - R_{\mathrm{mix}}}{L_{\mathrm{AX}}} \right) \tag{4.39}$$

This parameter represents the characteristic angle at which the swirling flow would reach the chamber wall at an axial distance equal to L_{AX}. Three distinct geometric configurations arise depending on the relative magnitudes of γ and α, illustrated in Figure 4.23.

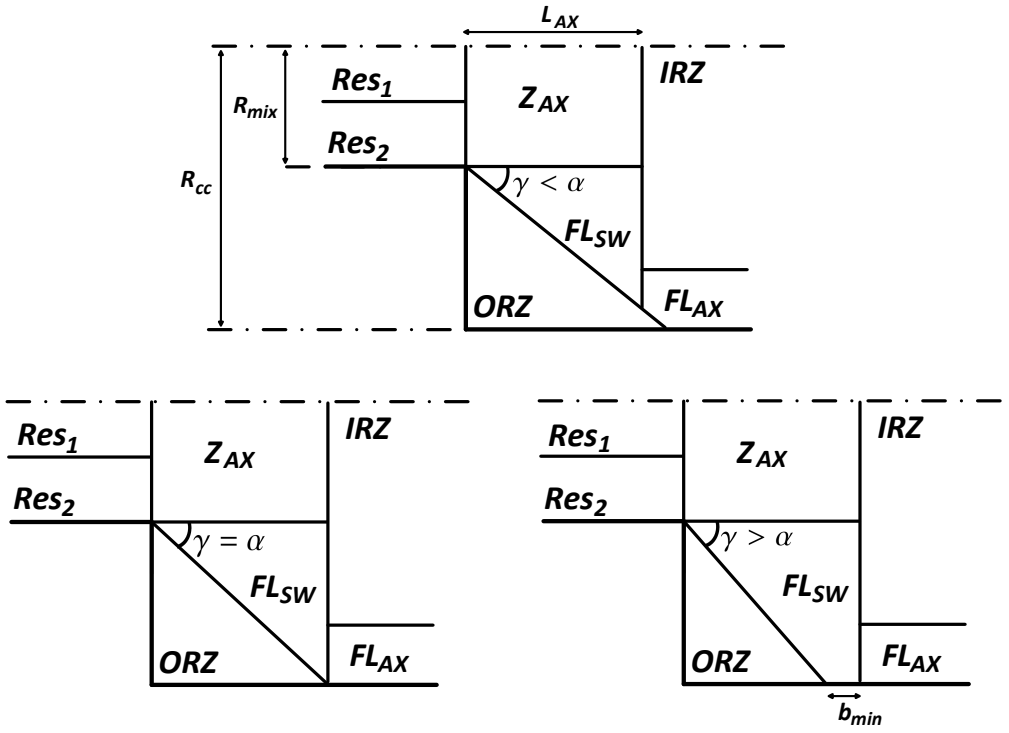

Fig. 4.23 Schematic representation of the volume allocation configurations for the swirling flame zone ($\mathrm{FL_{SW}}$) depending on the relationship between the deflection angle γ and the geometric parameter α. Top: $\gamma < \alpha$, the swirling flow remains confined within a conical region. Bottom left: $\gamma = \alpha$, the swirling flow contacts the chamber wall at L_{AX}. Bottom right: $\gamma > \alpha$, the swirling flow impinges on the wall upstream of L_{AX}, with b_{min} denoting the axial contact location. The volume formulations adapt to these geometric configurations to maintain physical consistency.

When $\gamma = \alpha$, the swirling flame zone contacts the chamber wall precisely at the axial distance L_{AX}. When $\gamma < \alpha$, the swirling flow remains confined within a conical region that does not reach the chamber wall. When $\gamma > \alpha$, the swirling flow impinges on the wall upstream of L_{AX}, yielding a trapezoidal profile. The corresponding volumes are:

$$V_{\mathrm{FL_{SW}}} = \begin{cases} \pi L_{\mathrm{AX}}(R_{\mathrm{cc}} - R_{\mathrm{mix}}) \left(R_{\mathrm{mix}} + \dfrac{R_{\mathrm{cc}} - R_{\mathrm{mix}}}{3} \right) & \gamma = \alpha \\ \pi L_{\mathrm{AX}}^2 \tan(\gamma) \left(R_{\mathrm{mix}} + \dfrac{L_{\mathrm{AX}} \tan(\gamma)}{3} \right) & \gamma < \alpha \\ \pi (R_{\mathrm{cc}} - R_{\mathrm{mix}}) R_{\mathrm{mix}} (b_{\mathrm{min}} + L_{\mathrm{AX}}) + \dfrac{\pi (R_{\mathrm{cc}} - R_{\mathrm{mix}})^2}{3} (L_{\mathrm{AX}} + 2b_{\mathrm{min}}) & \gamma > \alpha \end{cases} \tag{4.40}$$

where $b_{\mathrm{min}} = L_{\mathrm{AX}} - (R_{\mathrm{cc}} - R_{\mathrm{mix}}) \cot(\gamma)$ denotes the axial coordinate at which the swirling flow first contacts the chamber wall.

The $V_{\mathrm{FL_{AX}}}$ and V_{OUT} volumes were treated as optimization parameters and expressed as a fraction of the total chamber volume:

$$V_{\mathrm{FL_{AX}}} = \theta_9 V_{\mathrm{cc}} \tag{4.41}$$

$$V_{\mathrm{OUT}} = \theta_{10} V_{\mathrm{cc}} \tag{4.42}$$

where θ_9 and θ_{10} were determined through the optimization procedure alongside the other network parameters.

The inner recirculation zone volume was calculated as the difference between the total combustion chamber volume and the sum of all other defined zone volumes:

$$V_{\mathrm{IRZ}} = V_{\mathrm{cc}} - V_{\mathrm{Z_{AX}}} - V_{\mathrm{FL_{SW}}} - V_{\mathrm{ORZ}} - V_{\mathrm{FL_{AX}}} - V_{OUT} \tag{4.43}$$

Since the swirl number S depends on operating conditions through the density ratio and mass flow distribution (Equation 4.35), and since the deflection angle γ is directly related to S, this formulation enables dynamic volume recalculation throughout the simulations to account for variations in inlet temperature, equivalence ratio, and air mass flow rate. Furthermore, the optimization of network parameters compensates for potential systematic biases introduced by the simplified volume allocation strategy.

Parametric Formulation and Mass Flow Allocation

The determination of mass flow distributions between reactors requires specification of numerous flow splits that cannot be uniquely determined from continuity equations alone. To render the system tractable for optimization while maintaining physical interpretability, ten dimensionless parameters were defined to characterize the network configuration.

The mixing tube representation through two distinct inlet reservoirs necessitated the determination of individual stream properties from known global combustor operating conditions. Given the total air mass flow rate $\dot{m}_{\text{air}}$ and the global equivalence ratio ϕ, the individual equivalence ratios ϕ_1, ϕ_2 and mass flow rates $\dot{m}_{\text{air},1}$, $\dot{m}_{\text{air},2}$ for each reservoir stream required explicit calculation.

The parametric system was formulated through three dimensionless splitting parameters:

$$\theta_1 = \frac{\phi_1}{\phi_2} \tag{4.44}$$

$$\theta_2 = \frac{\dot{m}_{\text{air},1}}{\dot{m}_{\text{air},2}} \tag{4.45}$$

$$\theta_3 = \frac{\dot{m}_{2,\text{ORZ}}}{\dot{m}_2} \tag{4.46}$$

where θ_1 establishes the relative stoichiometry of the two streams, θ_2 controls the mass flow partitioning between the reservoirs, and θ_3 determines the fraction of mass flow from Res_2 directed to the outer recirculation zone. The complete system of equations governing the mixing tube is:

$$\begin{cases} \dot{m}_{\text{fuel}} = f_{\text{st}}(\phi_1 \dot{m}_{\text{air},1} + \phi_2 \dot{m}_{\text{air},2}) \\ \dot{m}_{\text{air}} = \dot{m}_{\text{air},1} + \dot{m}_{\text{air},2} \\ \dot{m}_2 = \dot{m}_{2,\text{FL}_{\text{SW}}} + \dot{m}_{2,\text{ORZ}} \\ \theta_1 = \dfrac{\phi_1}{\phi_2} \\ \theta_2 = \dfrac{\dot{m}_{\text{air},1}}{\dot{m}_{\text{air},2}} \\ \theta_3 = \dfrac{\dot{m}_{2,\text{ORZ}}}{\dot{m}_2} \end{cases} \tag{4.47}$$

where $f_{\text{st}} = 0.0292$ is the stoichiometric fuel-to-air mass ratio for hydrogen-air mixture, $\dot{m}_{\text{fuel}}$ is the total fuel mass flow rate, and $\dot{m}_{2,\text{FL}_{\text{SW}}}$ and $\dot{m}_{2,\text{ORZ}}$ denote the portions of mass flow from Res_2 directed to the swirling flame zone and outer recirculation zone, respectively.

Five additional parameters govern the mass flow distribution at critical bifurcation points within the combustion chamber:

$$\theta_4 = \frac{\dot{m}_B}{\dot{m}_{2,\text{FL}_{\text{SW}}} + \dot{m}_C + \dot{m}_E + \dot{m}_A} \tag{4.48a}$$

$$\theta_5 = \frac{\dot{m}_G}{T_{\text{FH}}} \tag{4.48b}$$

$$\theta_6 = \frac{\dot{m}_E}{T_{\text{FH}}(1-\theta_5)} \tag{4.48c}$$

$$\theta_7 = \frac{\dot{m}_I}{T_{\text{FH}}(1-\theta_5)(1-\theta_6)} \tag{4.48d}$$

$$\theta_8 = \frac{\dot{m}_H}{\dot{m}_D + \dot{m}_I} \tag{4.48e}$$

where $T_{\text{FH}} = \dot{m}_F + \dot{m}_H$ represents an intermediate mass flow variable. These parameters have the following physical interpretations:

- θ_4: Controls the split of the swirling flame zone outflow between the outer recirculation zone and the axial flame zone, representing the fraction directed to ORZ.

- θ_5, θ_6, θ_7: Prescribe a hierarchical redistribution of the IRZ outflow among its four exit pathways. Specifically, θ_5 defines the fraction returning to Z_{AX}; θ_6 determines the fraction of the remaining portion entering FL_{SW}; θ_7 specifies the fraction of the subsequent remainder proceeding to FL_{AX}. The residual fraction constitutes the IRZ discharge to the chamber exit.
- θ_8: Quantifies the fraction of FL_{AX} outflow that recirculates to the inner recirculation zone, with the complement exiting the main chamber.

This parameterization ensures that all splitting fractions remain strictly positive and bounded within the interval $[0,1]$, while employing the minimum number of independent parameters necessary to fully specify the network topology.

The complete optimization parameter vector is therefore:

$$\vec{\Theta} = [\theta_1, \theta_2, \theta_3, \theta_4, \theta_5, \theta_6, \theta_7, \theta_8, \theta_9, \theta_{10}] \tag{4.49}$$

From the mass balance equations and parameter definitions, a relation for T_{FH} as a function of $\dot{m}_A$ can be obtained by algebraic manipulation:

$$T_{FH} = \frac{\dot{m}_A(1-\theta_8) - \dot{m}_1 - \theta_8(\dot{m}_{2,FL_{SW}} + \dot{m}_{2,ORZ})}{\mathscr{D}(\vec{\Theta})} \tag{4.50}$$

where $\mathscr{D}(\vec{\Theta})$ represents a denominator function that depends on the parameter vector:

$$\mathscr{D}(\vec{\Theta}) = \theta_5\theta_6\theta_7\theta_8 - \theta_5\theta_6\theta_8 - \theta_5\theta_7\theta_8 + \theta_5 - \theta_6\theta_7\theta_8 + \theta_6\theta_8 + \theta_7\theta_8 - 1 \tag{4.51}$$

With T_{FH} expressed as a function of $\dot{m}_A$, all remaining mass flow rates can be computed sequentially. Additionally, a control mechanism enforces that the equivalence ratio in the swirling flame zone equals the global equivalence ratio:

$$f_{\mathrm{st,H_2/O_2}}\phi = \frac{\dot{m}_{2,\mathrm{FL_{SW}}} Z_{H,\mathrm{Res_2}} + \dot{m}_A Z_{H,\mathrm{Z_{AX}}} + \dot{m}_E Z_{H,\mathrm{IRZ}} + \dot{m}_B Z_{H,\mathrm{ORZ}} - \dot{m}_C Z_{H,\mathrm{FL_{SW}}} - \dot{m}_D Z_{H,\mathrm{FL_{SW}}}}{\dot{m}_{2,\mathrm{FL_{SW}}} Z_{O,\mathrm{Res_2}} + \dot{m}_A Z_{O,\mathrm{Z_{AX}}} + \dot{m}_E Z_{O,\mathrm{IRZ}} + \dot{m}_B Z_{O,\mathrm{ORZ}} - \dot{m}_C Z_{O,\mathrm{FL_{SW}}} - \dot{m}_D Z_{O,\mathrm{FL_{SW}}}} \tag{4.52}$$

where $Z_{H,i}$ and $Z_{O,i}$ represent the elemental mass fractions of hydrogen and oxygen in reactor i, and $f_{\mathrm{st,H_2/O_2}} = 0.125$ is the stoichiometric mass ratio of hydrogen to oxygen.

The mass flow rate $\dot{m}_A$ is determined by solving a nonlinear residual equation derived from Equations (4.52) and (4.50), employing a hybrid bracketing algorithm combining grid search and bisection. During the transient simulation, $\dot{m}_A$ and T_{FH} are dynamically recomputed at regular intervals to account for the temporal evolution of species compositions within each reactor.

Optimization Procedure

The ten parameters constituting $\vec{\Theta}$ were calibrated by minimizing the discrepancy between CRN predictions and experimental $\mathrm{NO_x}$ measurements across multiple equivalence ratios at the reference operating condition. Several optimization algorithms were evaluated for this task [89], including Particle Swarm Optimization [73], Differential Evolution [146], Bayesian Optimization [66, 133], ADAptive Moment estimation [77], and a hybrid Differential Evolution plus L-BFGS-B strategy [11]. Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [57, 58] proved to be the most suitable algorithm, achieving the lowest median score and the highest success rate across twenty independent runs, and is therefore adopted here [89].

The objective function was formulated to quantify the weighted relative error:

$$f = \sqrt{\sum_{i=1}^{n} w_i \left(\frac{\mathrm{NO}_{x,\mathrm{sim},i} - \mathrm{NO}_{x,\mathrm{exp},i}}{\mathrm{NO}_{x,\mathrm{exp},i}} \right)^2} \tag{4.53}$$

where $n = 4$ denotes the number of calibration points. To emphasize accuracy at conditions with higher absolute emission levels, magnitude-dependent weights were assigned:

$$w_i = \frac{\mathrm{NO}_{x,\mathrm{exp},i}^{0.85}}{\sum_{j=1}^{n} \mathrm{NO}_{x,\mathrm{exp},j}^{0.85}} \tag{4.54}$$

The exponent value of 0.85 was empirically selected to provide moderate emphasis on high-emission conditions without neglecting low-emission cases.

The calibration was performed using four equivalence ratio points ($\phi = 0.40$, 0.60, 0.87, and 0.96) at the reference operating condition ($\dot{m}_{\mathrm{air}} = 130$ kg/h, $T_{\mathrm{air}} = 623$ K). These points were selected to span the full range of equivalence ratios investigated, including both boundary values and intermediate conditions, ensuring representative coverage of the combustor operating envelope.

CMA-ES [57, 58] adapts the full covariance matrix of a multivariate normal sampling distribution, effectively learning both the magnitude and orientation of the search ellipsoid from the ranking of candidate solutions. This adaptation renders the algorithm invariant to orthogonal transformations of the search space, a property that is particularly advantageous for problems with non-separable parameter interactions [64]. The offspring population size is set to $\lambda = 11$, and the initial step size is set to $\sigma_0 = 0.5$ in normalized coordinates. Bound constraints are handled through a sigmoid transformation, and infeasible evaluations are assigned a large finite penalty rather than infinity to avoid corrupting the covariance matrix update.

The optimization parameters were constrained within physically meaningful bounds, listed together with the optimized parameter set in Table 4.10.

Table 4.10 Optimized CRN parameters obtained from CMA-ES calibration for the Naik and CRECK mechanisms at the reference operating condition ($\dot{m}_{\text{air}} = 130$ kg/h, $T_{\text{air}} = 623$ K). All optimizations targeted four equivalence ratio points: $\phi = 0.40$, 0.60, 0.87, and 0.96.

Parameter	Naik	CRECK	Range	Definition
θ_1	0.223	0.211	[0.20, 5.00]	ϕ_1/ϕ_2
θ_2	0.201	0.197	[0.20, 5.00]	$\dot{m}_{\text{air},1}/\dot{m}_{\text{air},2}$
θ_3	0.286	0.300	[0.00, 1.00]	$\dot{m}_{2,\text{ORZ}}/\dot{m}_2$
θ_4	0.467	0.444	[0.00, 1.00]	$\dot{m}_B/(\dot{m}_{2,\text{FL}_{\text{SW}}} + \dot{m}_C + \dot{m}_E + \dot{m}_A)$
θ_5	0.056	0.054	[0.00, 1.00]	$\dot{m}_G/T_{\text{FH}}$
θ_6	0.931	0.885	[0.00, 1.00]	$\dot{m}_E/[T_{\text{FH}}(1-\theta_5)]$
θ_7	0.555	0.566	[0.00, 1.00]	$\dot{m}_I/[T_{\text{FH}}(1-\theta_5)(1-\theta_6)]$
θ_8	0.362	0.363	[0.00, 1.00]	$\dot{m}_H/(\dot{m}_D + \dot{m}_I)$
θ_9	0.051	0.053	[0.05, 0.15]	$V_{\text{FL}_{\text{AX}}}/V_{\text{cc}}$
θ_{10}	0.202	0.210	[0.20, 0.45]	$V_{\text{OUT}}/V_{\text{cc}}$

The CRN was first calibrated using the Naik mechanism, and the resulting optimal parameter vector $\vec{\Theta}^*_{\text{Naik}}$ was subsequently used as the starting point for a local CMA-ES optimization with the CRECK mechanism, exploiting the proximity of the two solutions in parameter space. Since the CRECK optimization must be close to the Naik optimization, i.e. below 5%, it has been allowed to fall outside the prescribed bounds in Table 4.10 for the parameters θ_1, θ_2, θ_9 and θ_{10}. The parameter θ_6, controlling the fraction of IRZ outflow entering FL_{SW}, shows the largest variation (0.931 for Naik versus 0.885 for CRECK), reflecting different residence time requirements in the swirling flame zone to achieve equivalent NO_x predictions with the two chemical kinetics schemes. The computational cost of the optimization procedure was approximately 40 CPU-hours on average, achieving a global best of $J = 0.039$ against an imposed requirement of 0.05, a cost negligible compared to equivalent parametric CFD studies.

The distribution of the best-found parameter vectors across twenty independent optimization runs, shown in Fig. 4.24, reveals a clear distinction between well-identified and poorly-identified parameters. Several parameters are tightly constrained across all seeds. The inlet mass-flow ratio $\theta_2 = \dot{m}_{\text{air},1}/\dot{m}_{\text{air},2}$ converges consistently to $\theta_2 \approx 0.20$, indicating that nearly all of the air is routed through the second reservoir regardless of the initial conditions. The volume fractions θ_9 ($V_{FL_{AX}}/V_{CC}$)

and θ_{10} (V_{OUT}/V_{CC}) similarly cluster near their lower bounds ($\theta_9 \approx 0.05$, $\theta_{10} \approx 0.20$), demonstrating that the best CRN configurations assign the minimum permitted volume to $\mathrm{FL_{AX}}$ and OUT, leaving the maximum residual volume to the inner recirculation zone. The IRZ-to-$\mathrm{FL_{SW}}$ recirculation fraction θ_6 converges to high values ($\theta_6 \approx 0.93$), reflecting the physical dominance of the recirculation feedback from the IRZ to the swirling flame zone. In contrast, the ORZ fraction (θ_3), the $\mathrm{FL_{SW}}$-to-ORZ split (θ_4), the IRZ-to-Z_{AX} fraction (θ_5), and the IRZ-to-$\mathrm{FL_{AX}}$ fraction (θ_7) exhibit substantial variability across seeds. This dispersion suggests that multiple distinct configurations produce comparably low objective values, reflecting compensating effects among the internal mass-flow redistribution pathways. This observation carries direct implications for the uncertainty quantification analysis presented in Sec. 4.3.7: the parameters θ_2, θ_6, θ_9, and θ_{10} must be determined precisely, whereas the CRN predictions are comparatively insensitive to perturbations in θ_3, θ_4, θ_5, and θ_7.

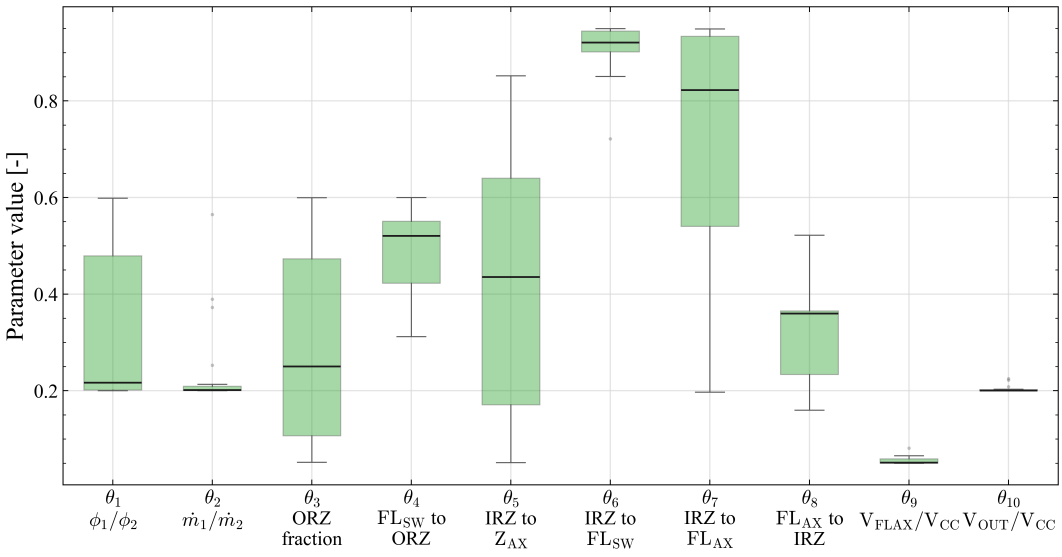

Fig. 4.24 Distribution of best-found parameter vectors across 20 independent seeds.

4.3.5 Results and Validation

Calibration Point Performance

The optimized CRN models were first evaluated at the four calibration points to assess the quality of the optimization. Figure 4.25 presents the comparison between

experimental NO_x emissions and CRN predictions using the Naik and CRECK mechanisms at the reference operating condition.

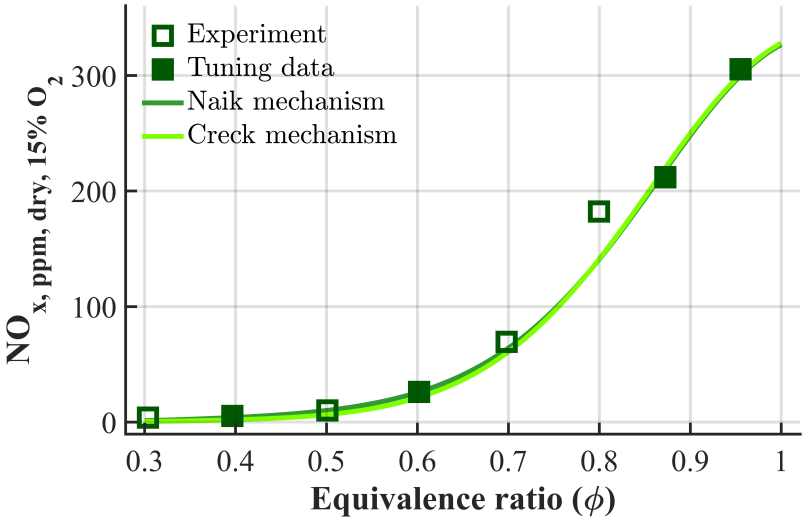

Fig. 4.25 Comparison of experimental NO_x emissions with CRN predictions using optimized parameters for the Naik and CRECK mechanisms at the reference operating condition ($\dot{m}_{air} = 130$ kg/h, $T_{air} = 623$ K). Filled squares indicate experimental data points used for CRN optimization ($\phi = 0.40$, 0.60, 0.87, 0.96), while open squares represent validation points not included in the optimization procedure. Both mechanisms yield nearly identical predictions across the full equivalence ratio range.

The two mechanisms yield nearly identical predictions across the entire equivalence ratio range, with a mean absolute difference of approximately 2.5 ppm. This close agreement indicates that the optimized CRN parameters effectively compensate for differences in the underlying chemical kinetics, producing consistent emission estimates regardless of the selected mechanism.

At the optimization points, the mechanisms demonstrate satisfactory agreement with experimental data. The highest relative errors occur at $\phi = 0.40$, where the Naik and CRECK mechanisms predict 4.01 and 2.00 ppm respectively, compared to the experimental value of 5.28 ppm. At higher equivalence ratios, the agreement improves substantially: at $\phi = 0.87$ and $\phi = 0.96$, both mechanisms reproduce experimental values within 3% and 4% relative error, respectively.

The validation points, which were excluded from the optimization procedure, provide an independent assessment of model predictive capability. For the majority of validation points ($\phi = 0.30$, 0.50, and 0.70), the Naik and CRECK mechanisms yield absolute errors below 7 and 10 ppm, respectively, demonstrating reasonable

interpolation performance for intermediate equivalence ratios. A notable discrepancy occurs at $\phi = 0.80$, where both mechanisms predict approximately 141 ppm against an experimental value of 182.3 ppm, resulting in an absolute error of approximately 41 ppm. However, this experimental point appears to deviate from the trend exhibited by the remaining data, suggesting potential measurement anomalies or localized combustion instabilities during acquisition. Excluding this point, the mean absolute error across validation points decreases substantially.

Quantitative comparison of experimental and simulated NO_x emissions at the reference condition is presented in Table 4.11.

Table 4.11 Comparison of experimental and simulated NO_x emissions for the reference operating condition ($\dot{m}_{air} = 130$ kg/h, $T_{air} = 623$ K). Absolute and relative errors are defined as $\varepsilon_{abs} = \text{Sim} - \text{Exp}$ and $\varepsilon_{rel} = (\text{Sim} - \text{Exp})/\text{Exp} \times 100$ [%], respectively; negative values indicate under-prediction by the CRN, positive values indicate over-prediction. This sign convention is adopted in all subsequent comparison tables.

		NO_x [ppm, dry, 15% O_2]			ε_{abs} [ppm]		ε_{rel} [%]	
Type	ϕ	Exp.	Naik	CRECK	Naik	CRECK	Naik	CRECK
Tuning	0.40	5.28	4.01	2.00	−1.28	−3.28	−24.2	−62.0
	0.60	26.13	26.08	21.47	−0.05	−4.66	−0.2	−17.8
	0.87	212.07	218.24	220.04	+6.17	+7.97	+2.9	+3.8
	0.96	305.54	299.86	301.10	−5.68	−4.44	−1.9	−1.4
Validation	0.30	3.93	1.37	0.61	−2.55	−3.31	−65.0	−84.4
	0.50	10.12	10.16	6.52	+0.05	−3.60	+0.5	−35.6
	0.70	69.40	63.08	59.77	−6.33	−9.63	−9.1	−13.9
	0.80	182.33	140.68	141.04	−41.65	−41.30	−22.8	−22.6

The steady-state reactor conditions for a representative case are illustrated in Figure 4.26, which depicts the temperature, pressure, equivalence ratio, and major species composition in each reactor using the Naik mechanism at the reference condition with $\phi = 0.60$. The control mechanism successfully enforces that $\phi_{FL_{SW}}$ is close to the global equivalence ratio equal to 0.6.

Validation Across Operating Conditions

Having calibrated the CRN at the reference condition, the crucial test of predictive capability involves applying the same network structure and parameters to the full experimental matrix without recalibration. The validated CRN was evaluated across

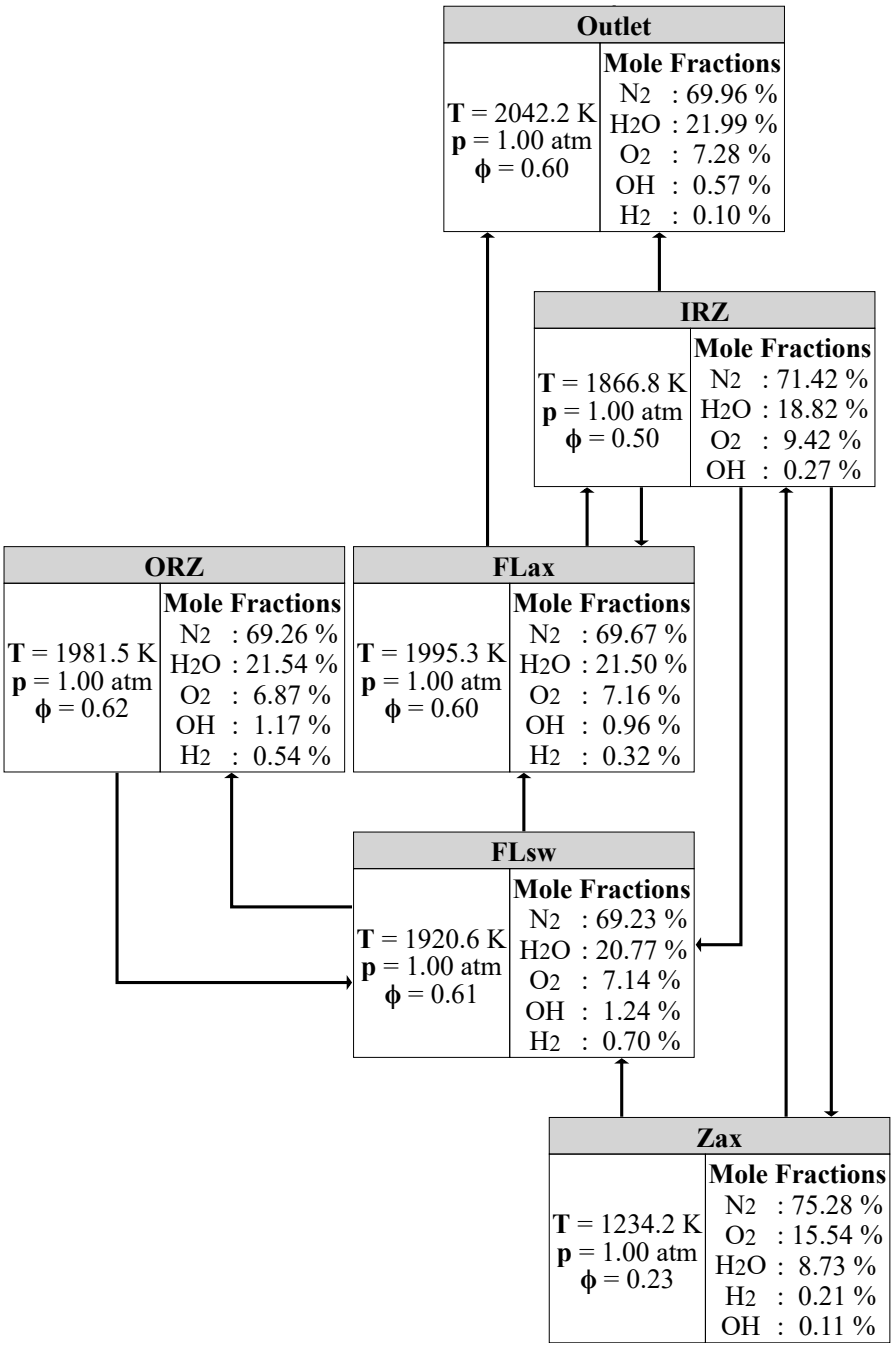

Fig. 4.26 Chemical reactor network topology at the reference operating condition ($\dot{m}_{\text{air}} = 130$ kg/h, $T_{\text{air}} = 623$ K) and $\phi = 0.60$. Each reactor displays the steady-state temperature, pressure, local equivalence ratio, and major species mole fractions. The arrows indicate mass flow connections between zones.

inlet air temperatures from 313 K to 693 K, air mass flow rates from 116 kg/h to 255 kg/h, and equivalence ratios from 0.30 to 0.96.

Figures 4.27, 4.28, and 4.29 present comparisons between experimental NO_x measurements and CRN predictions using both mechanisms for the three non-reference operating conditions.

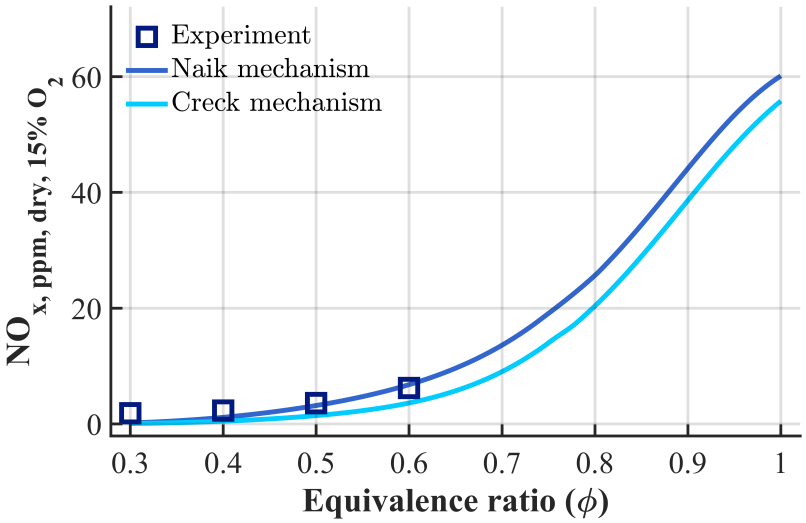

Fig. 4.27 Comparison of experimental and CRN-predicted NO_x emissions at operating condition 1: $\dot{m}_{air} = 255$ kg/h, $T_{air} = 313$ K. Symbols represent experimental measurements; lines indicate CRN predictions using the Naik and CRECK mechanisms with parameters from Table 4.10 (no recalibration).

Consistent with the reference condition behaviour, both chemical mechanisms exhibit a tendency to underestimate NO_x emissions at low equivalence ratios across all operating conditions. This systematic bias suggests that the simplified reactor network may not fully capture the complex interaction between turbulent mixing and chemistry in the ultra-lean regime, where hotspots, local extinction and re-ignition phenomena become significant. Notably, the CRECK mechanism demonstrates a more pronounced underestimation than the Naik mechanism in this lean combustion regime, as clearly illustrated in Figure 4.27, where the reduced emission levels allow detailed observation of the discrepancy.

Despite this tendency, both mechanisms maintain reasonable accuracy within the experimentally validated range. Specifically, the maximum absolute error across all non-reference operating conditions does not exceed approximately 3.0 ppm for the Naik mechanism and 7 ppm for the CRECK mechanism. Quantitative comparison

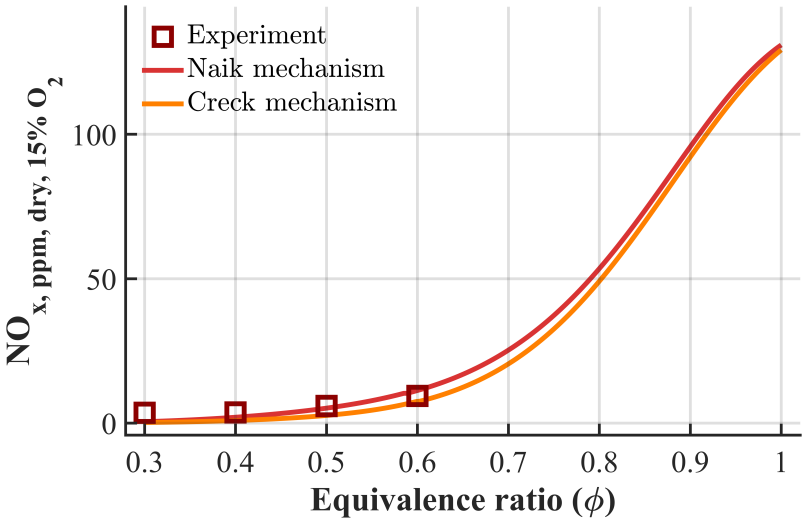

Fig. 4.28 Comparison of experimental and CRN-predicted NO_x emissions at operating condition 2: $\dot{m}_{air} = 180$ kg/h, $T_{air} = 453$ K. Symbols represent experimental measurements; lines indicate CRN predictions using the Naik and CRECK mechanisms with parameters from Table 4.10 (no recalibration).

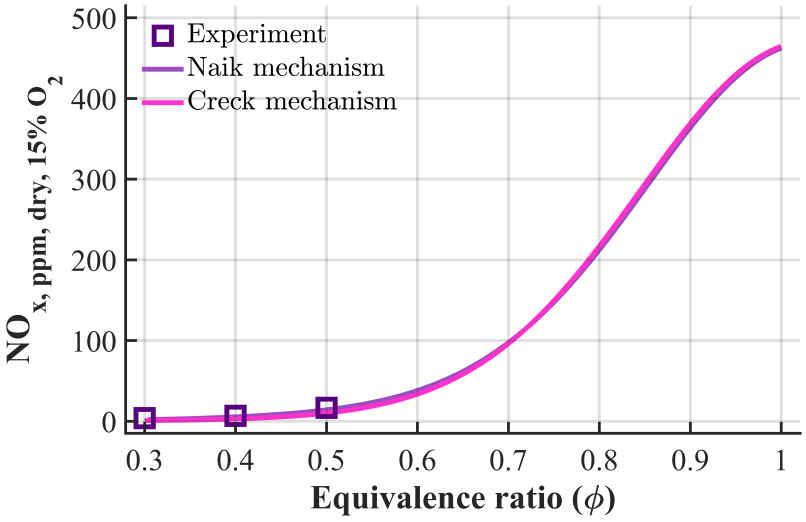

Fig. 4.29 Comparison of experimental and CRN-predicted NO_x emissions at operating condition 4: $\dot{m}_{air} = 116$ kg/h, $T_{air} = 693$ K. Symbols represent experimental measurements; lines indicate CRN predictions using the Naik and CRECK mechanisms with parameters from Table 4.10 (no recalibration).

between experimental measurements and CRN predictions for each non-reference operating condition is presented in Table 4.12.

Table 4.12 Comparison of experimental and simulated NO_x emissions for the non-reference operating conditions using the Naik and CRECK chemical mechanisms. All emissions are reported in ppm at 15% O_2, dry basis. The same optimised parameters from Table 4.10 were used without recalibration. Sign convention as defined in Table 4.11.

Mech.	ϕ	$\dot{m}_{air} = 255$ kg/h, $T_{air} = 313$ K				$\dot{m}_{air} = 180$ kg/h, $T_{air} = 453$ K			
		Exp. [ppm]	Sim. [ppm]	ε_{abs} [ppm]	ε_{rel} [%]	Exp. [ppm]	Sim. [ppm]	ε_{abs} [ppm]	ε_{rel} [%]
Naik	0.30	1.83	0.18	−1.65	−90.4	3.56	0.50	−3.06	−85.9
	0.40	2.32	1.16	−1.16	−50.0	3.69	2.07	−1.62	−44.0
	0.50	3.61	3.24	−0.37	−10.3	5.94	5.27	−0.67	−11.3
	0.60	6.22	6.91	+0.69	+11.1	9.39	11.53	+2.14	+22.8
CRECK	0.30	1.83	0.07	−1.76	−96.1	3.56	0.21	−3.35	−94.2
	0.40	2.32	0.48	−1.84	−79.2	3.69	0.92	−2.77	−75.1
	0.50	3.61	1.47	−2.14	−59.3	5.94	2.71	−3.23	−54.3
	0.60	6.22	3.70	−2.52	−40.4	9.39	7.57	−1.82	−19.3

Table 4.13 Continuation of Table 4.12: comparison for operating condition 4 ($\dot{m}_{air} = 116$ kg/h, $T_{air} = 693$ K).

Mechanism	ϕ	$\dot{m}_{air} = 116$ kg/h, $T_{air} = 693$ K			
		Exp. [ppm]	Sim. [ppm]	ε_{abs} [ppm]	ε_{rel} [%]
Naik	0.30	3.90	1.94	−1.95	−50.2
	0.40	6.75	5.54	−1.21	−17.9
	0.50	16.64	14.48	−2.16	−13.0
CRECK	0.30	3.90	0.90	−2.99	−76.8
	0.40	6.75	3.00	−3.74	−55.5
	0.50	16.64	10.29	−6.35	−38.1

Figure 4.30 shows the CRN predictions across all four operating conditions using the Naik mechanism, illustrating the model's capability to capture emission trends as functions of both equivalence ratio or combustor exit temperature across the full operating envelope. In particular, looking at NO_x emissions respect to the flame temperature in the outlet reactor, the CRN captures the exponential formation of NO_x after approximately 1800 K.

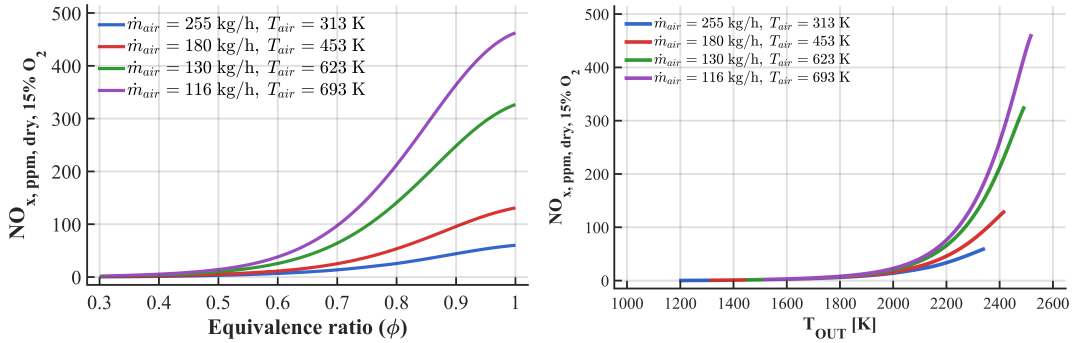

(a) NO_x emissions vs. equivalence ratio (b) NO_x emissions vs. exit temperature

Fig. 4.30 CRN predictions using the Naik mechanism across all four operating conditions. (a) NO_x emissions increase monotonically with equivalence ratio. (b) NO_x emissions exhibit the expected exponential dependence on combustor temperature and diminish as the air mass flow rate increases, corresponding to shorter residence times.

The influence of inlet air temperature on NO_x formation is evident when comparing operating conditions at fixed equivalence ratios. At $\phi = 0.60$, predicted NO_x emissions increase from 6.91 ppm at the lowest inlet temperature condition ($T_{\text{air}} = 313$ K) to 39.11 ppm at the highest ($T_{\text{air}} = 693$ K, extrapolated beyond available experimental data). This sevenfold increase reflects the combined effects of elevated combustor temperatures and modified reaction kinetics at higher preheat levels.

The effect of air mass flow rate on NO_x emissions is demonstrated by comparing operating conditions at similar exit temperatures. At $T_{\text{exit}} \approx 2300$ K, the highest mass flow rate condition ($\dot{m}_{\text{air}} = 255$ kg/h) yields approximately 50 ppm, whereas the lowest mass flow rate ($\dot{m}_{\text{air}} = 116$ kg/h) produces nearly 145 ppm. This inverse relationship between mass flow rate and NO_x emissions arises from extended residence times in high-temperature reaction zones at lower flow rates, providing additional time for thermal NO_x formation.

Physical Insights and Limitations

Several limitations of the CRN approach merit discussion. First, the network was calibrated using experimental data at only one operating condition (130 kg/h, 623 K), though across multiple equivalence ratios. The demonstrated transferability to substantially different inlet temperatures (313 K to 693 K) and mass flow rates (116 kg/h to 255 kg/h) provides confidence in the physical basis of the network, but

additional calibration data at operating conditions even further from the reference one would strengthen the validation.

Second, the systematic underestimation of NO_x at very lean conditions ($\phi < 0.4$) suggests that the PSR approximation may not adequately represent the flame structure in this regime. Ultra-lean flames exhibit increased susceptibility to local extinction and partial reaction, phenomena that cannot be captured by well-stirred reactor models. For applications requiring accurate predictions at very lean conditions, more sophisticated submodels or additional reactor zones may be necessary.

Third, the volume allocation strategy assumes that the flow topology remains sufficiently similar across operating conditions to justify using fixed geometric relations. While the demonstrated predictive capability validates this assumption within the tested envelope, dramatic changes in flow regime (e.g., transition from swirl-dominated to jet-dominated flow) could invalidate the volume allocations and require network reconfiguration.

4.3.6 Multi-Fidelity Comparative Assessment

The Chemical Reactor Network results presented in Section 4.3.5 are contextualised here within a broader multi-fidelity modelling hierarchy for the AHEAD combustor, whose computational fluid dynamics component is described in Section 3.3.5. The hierarchy encompasses three fidelity levels: the Chemical Reactor Network with both the Naik and CRECK mechanisms (CRN); two-dimensional axisymmetric RANS with both mechanisms (RANS–2D); and three-dimensional RANS with Explicit Algebraic Reynolds Stress Model (EARSM) turbulence closure [155, 156, 6] and the Naik mechanism (RANS–3D EARSM). The three-dimensional formulation was evaluated exclusively with the Naik mechanism owing to its substantially higher computational cost. All CFD simulations were performed on Intel® Xeon® Gold 6338 CPUs at 2.00 GHz using Ansys® Fluent 2024 R1.

To ensure numerical consistency across both modelling dimensions, the two-dimensional axisymmetric and three-dimensional simulations utilize identical discretization frameworks featuring second-order upwind schemes, a coupled pressure-velocity solver, and the PRESTO! interpolation method for pressure fields [6]. The RANS–2D simulations employ the same two-dimensional axisymmetric mesh described in Section 3.3.4, comprising approximately 2×10^5 poly-hexcore cells

(Fig. 3.7). The RANS–3D EARSM simulations, by contrast, require a dedicated three-dimensional mesh, described below.

Three-Dimensional Mesh

A poly-hexcore meshing strategy was evaluated for the three-dimensional simulations using two grids of 2.1×10^6 and 3.2×10^6 cells. On the coarser grid, the target volume element size is 2 mm within the flame, mixing tube, swirler, and injection plate, bounded between 0.5 mm and 4 mm depending on local geometric features. The surface mesh is 1 mm on the combustion chamber walls and 0.5 mm on the mixing tube and swirler, again bounded between 0.5 and 4 mm. The refined grid halves the target volume element size in the reactive regions to 1 mm and reduces the local volume maximum to 2 mm, while the surface specifications remain unchanged.

Comparison between the two refinement levels revealed discrepancies in the temperature field, and consequently in NO_x production, particularly within the outer recirculation zone. Therefore, to resolve these flow structures more accurately, all three-dimensional reactive flow simulations were performed using the refined 3.2×10^6 cell mesh, shown in Fig. 4.31. The coarser 2.1×10^6 cell solution was used solely to provide an interpolated initial condition for the refined-mesh runs, thereby reducing the number of iterations required to reach convergence; this procedure is reflected in the computational cost reported in Section 4.3.6.

NO_x Emission Database

The complete multi-fidelity numerical database is reported in Tables 4.14 and 4.15. Table 4.14 reports predictions at the eighteen operating points defined by constant $u_0 = 70$ m/s, where u_0 denotes the mean air axial velocity in the mixing tube in the absence of fuel, computed as:

$$u_0 = \frac{\dot{m}_{\text{air}}}{A_{\text{mix}}\,\rho_{\text{air}}} \tag{4.55}$$

These points correspond to the four reference inlet air temperatures $T_{\text{in}} = 313$, 453, 623, and 693 K, at baseline mass flow rates of 255, 180, 130, and 116 kg/h respectively. All emission values are in parts per million by volume on a dry

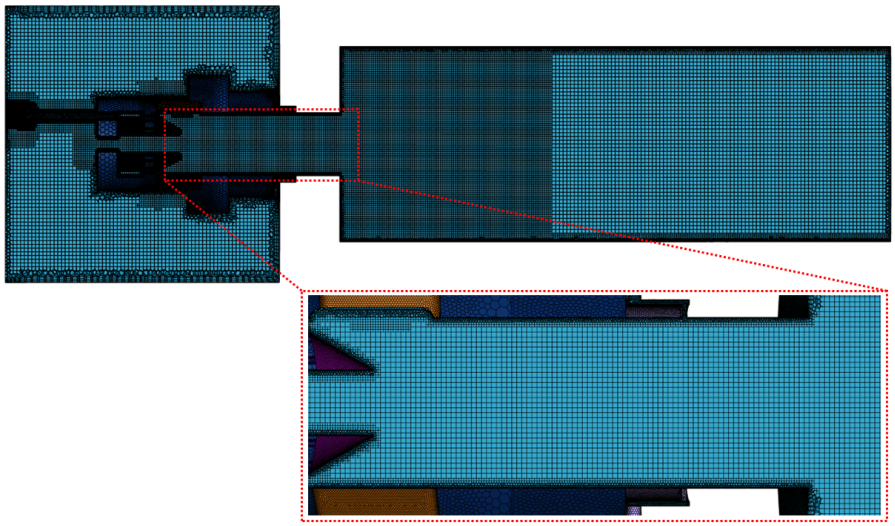

Fig. 4.31 Three-dimensional poly-hexcore mesh of the AHEAD combustor employed for the RANS–3D EARSM simulations (3.2×10^6 cells). The refined target element size of 1 mm is applied within the reactive regions, including the flame, mixing tube, swirler, and injection plate.

basis, corrected to 15% O_2. The RANS–3D EARSM simulation at $T_{\text{in}} = 313$ K, $\dot{m}_{\text{air}} = 255$ kg/h, and $\phi = 0.30$ did not achieve stable ignition; the corresponding entry is reported as "–". Table 4.15 reports parametric sweeps in which the inlet air mass flow rate is varied at fixed equivalence ratio ($\phi = 0.40$ and $\phi = 0.60$).

Validation and Split-Regime Error Analysis

The eighteen operating points of Table 4.14 for which experimental measurements are available span nearly two orders of magnitude, from approximately 1.8 to 305 ppm. To prevent the high-emission points from dominating the aggregated error metrics, the dataset is partitioned into a low-NO_x regime ($\phi \in [0.30, 0.50]$, twelve points, $NO_x \lesssim 17$ ppm) and a high-NO_x regime ($\phi > 0.50$, six points, up to ~ 305 ppm). The Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE) computed for each model in each regime are reported in Table 4.16.

The two regimes display an expected and complementary error structure. In the low-NO_x regime, the absolute error is bounded (MAE ≤ 3.4 ppm for every model) while the relative error is large (29–68%), a consequence of the small absolute emission levels. In the high-NO_x regime the absolute error grows by roughly an order

Table 4.14 NO_x emissions at operating points with constant $u_0 = 70$ m/s. Values are reported in ppm, dry, corrected to 15% O_2. The RANS–3D EARSM entry marked "–" corresponds to flame extinction at that condition and is assigned a value of 0 ppm in the computation of the error metrics reported in Table 4.16.

T_{in} [K]	$\dot{m}_{air}$ [kg/h]	ϕ [–]	Exp. [ppm]	RANS–2D Naik [ppm]	RANS–2D CRECK [ppm]	RANS–3D EARSM, Naik [ppm]	CRN Naik [ppm]	CRN CRECK [ppm]
313	255	0.30	1.83	0.38	0.42	–	0.18	0.07
313	255	0.40	2.32	1.70	0.63	0.86	1.16	0.48
313	255	0.50	3.61	3.75	1.71	2.27	3.24	1.47
313	255	0.60	6.22	7.62	5.90	4.79	6.91	3.70
313	255	0.70	–	11.40	7.71	9.53	13.88	9.27
453	180	0.30	3.56	0.87	0.33	0.39	0.50	0.21
453	180	0.40	3.69	2.60	0.66	1.48	2.07	0.92
453	180	0.50	5.94	5.29	2.42	3.42	5.27	2.71
453	180	0.60	9.39	10.26	6.55	7.60	11.53	7.57
453	180	0.80	–	44.06	39.46	50.52	54.76	50.43
453	180	1.00	–	84.47	89.26	143.87	131.85	130.31
623	130	0.30	3.93	1.87	0.73	1.03	1.37	0.61
623	130	0.40	5.28	4.90	2.30	2.77	4.01	2.00
623	130	0.50	10.12	9.79	5.63	6.33	10.16	6.51
623	130	0.60	26.13	21.70	16.21	16.80	26.08	21.47
623	130	0.70	69.40	52.57	46.05	49.96	63.08	59.77
623	130	0.80	182.33	127.61	127.35	146.32	140.68	141.04
623	130	0.87	212.07	195.04	199.12	233.18	218.25	220.04
623	130	0.96	305.54	258.15	266.33	333.04	299.87	301.11
623	130	1.00	–	268.77	282.15	338.93	328.41	329.90
693	116	0.30	3.90	2.54	1.02	1.43	1.94	0.90
693	116	0.40	6.75	6.51	3.23	3.67	5.54	3.00
693	116	0.50	16.64	13.36	8.45	8.66	14.48	10.29
693	116	0.60	–	31.51	25.76	25.86	39.11	34.67
693	116	0.80	–	201.47	204.77	232.25	217.41	222.02
693	116	1.00	–	400.18	416.14	472.49	464.08	466.26

Table 4.15 NO_x emissions for the inlet air mass flow rate parametric sweeps at fixed equivalence ratio $\phi = 0.40$ (upper block) and $\phi = 0.60$ (lower block). Values are reported in ppm, dry, corrected to 15% O_2.

T_{in}	$\dot{m}_{air}$	ϕ	RANS–2D Naik	RANS–2D CRECK	RANS–3D EARSM, Naik	CRN Naik	CRN CRECK
[K]	[kg/h]	[–]	[ppm]	[ppm]	[ppm]	[ppm]	[ppm]
313	130	0.40	3.09	1.47	1.01	1.12	0.48
313	184	0.40	2.18	0.91	0.92	1.14	0.48
313	220	0.40	1.88	0.74	0.90	1.16	0.48
313	295	0.40	1.56	0.58	0.84	1.18	0.49
453	80	0.40	6.44	4.27	1.60	2.16	1.05
453	130	0.40	3.69	1.63	1.55	2.09	0.96
453	150	0.40	3.14	1.33	1.52	2.08	0.94
453	205	0.40	2.34	0.95	1.47	2.07	0.91
453	230	0.40	2.17	0.87	1.47	2.07	0.90
623	93	0.40	7.23	4.06	2.83	4.34	2.39
623	110	0.40	5.83	2.96	2.80	4.35	2.25
623	148	0.40	4.31	1.92	2.75	4.15	2.06
693	82	0.40	10.23	6.23	3.75	6.04	3.46
693	100	0.40	7.72	3.80	3.69	5.75	3.19
693	133	0.40	5.67	2.62	3.62	5.41	2.87
313	130	0.60	13.52	11.10	5.28	8.07	4.86
313	184	0.60	10.78	7.74	4.90	7.44	4.22
313	220	0.60	8.88	5.83	4.75	7.15	3.93
313	295	0.60	6.49	3.60	4.53	6.73	3.53
453	80	0.60	25.30	22.55	9.44	15.76	12.03
453	130	0.60	15.35	10.04	8.07	12.99	9.08
453	150	0.60	12.27	8.53	7.76	12.32	8.39
453	205	0.60	9.77	5.57	7.24	11.08	7.28
453	230	0.60	9.02	5.38	7.09	10.68	6.85
623	93	0.60	31.13	26.30	20.39	32.04	27.80
623	110	0.60	26.20	20.55	18.39	28.72	24.31
623	148	0.60	18.48	13.75	15.82	24.60	19.86
693	82	0.60	47.96	42.64	33.88	49.53	45.95
693	100	0.60	37.50	31.53	28.81	43.26	39.16
693	133	0.60	26.86	21.94	23.73	36.03	31.25

Table 4.16 Split-regime validation metrics computed against the experimental NO_x data of Table 4.14. The RANS–3D EARSM simulation that did not achieve stable ignition is assigned 0 ppm in the error calculation.

	Low-NO_x ($\phi \leq 0.50$)		High-NO_x ($\phi > 0.50$)	
Model	MAE [ppm]	MAPE [%]	MAE [ppm]	MAPE [%]
RANS–2D, Naik	1.19	28.9	20.38	18.1
RANS–2D, CRECK	3.34	66.0	20.51	22.3
RANS–3D EARSM, Naik	2.94	58.9	16.66	20.6
CRN, Naik	1.48	38.5	8.96	10.1
CRN, CRECK	3.20	67.6	10.33	17.1

of magnitude (9–21 ppm) while the MAPE drops to 10–22%. All five formulations reproduce the near-exponential rise of NO_x with ϕ seen at $T_{in} = 623$ K, where the experimental measurements span from 3.93 ppm at $\phi = 0.30$ to 305.54 ppm at $\phi = 0.96$.

Across the full operating envelope, CRN Naik achieves the lowest MAE (8.96 ppm) and the lowest MAPE (10.1%) in the high-NO_x regime, benefiting in part from the inclusion of the high-ϕ calibration points in the optimization procedure. In the low-NO_x regime, however, RANS–2D Naik is the most accurate formulation in both absolute (MAE = 1.19 ppm) and relative (MAPE = 28.9%) terms. The large relative errors observed at lean conditions for all models are consistent with the systematic under-prediction discussed in Section 4.3.5: at very lean equivalence ratios, localised extinction and re-ignition phenomena generate NO_x levels that exceed those reproduced by homogeneous reactor or steady-state RANS descriptions alike.

Fidelity–Accuracy Trade-off

Holding the chemical mechanism fixed at Naik, the RANS–3D EARSM formulation does not consistently deliver superior predictive accuracy relative to the lower-fidelity models. In the low-NO_x regime its MAE = 2.94 ppm and MAPE = 58.9% are intermediate between RANS–2D Naik and the CRECK-based predictions; in the high-NO_x regime it improves on RANS–2D Naik in absolute terms (16.19 vs. 20.97 ppm) but is outperformed by CRN Naik on both metrics.

The three Naik-based predictions remain coherent in emission trend across the entire ϕ sweep, but cross over near $\phi = 0.80$–1.00. At $\phi = 0.96$ ($T_{\text{in}} = 623$ K), RANS–3D EARSM over-predicts the experimental measurement (333.04 vs. 305.54 ppm), CRN Naik provides the closest match (299.87 ppm, also benefiting from calibration at this point), and RANS–2D Naik under-predicts (258.15 ppm). The elevated errors in the low-ϕ regime are primarily attributable to the absence of a transient formulation in both CFD and CRN, which precludes the representation of thermoacoustic instabilities and the localised hot spots characteristic of lean hydrogen flames. In the high-ϕ regime, by contrast, NO_x production is governed by local equivalence ratio fluctuations through the highly non-linear temperature dependence of thermal NO_x; consequently, a more complete description of turbulent mixing, as provided by the three-dimensional EARSM closure, yields a more consistent treatment than the two-dimensional axisymmetric approximation, though the CRN calibration procedure partially compensates for these effects.

Computational Cost

The computational cost of each fidelity level is evaluated in terms of average CPU hours over the corresponding simulation dataset. Only the Naik chemical mechanism is considered for this comparison, so as to isolate the effect of model fidelity from that of the chemical mechanism. The average costs of the four levels are summarized in Table 4.17.

The CRN represents the lowest-cost level of the hierarchy, requiring approximately 18 s on a single core, corresponding to 5.0×10^{-3} CPU h per operating point. The 2D axisymmetric RANS simulations required on average 13.85 CPU h per case. The transition from RANS–2D to RANS–3D EARSM increases the average cost to 247.67 CPU h per case, approximately 17.9 times the cost of a 2D simulation. This increase arises from the larger cell count, the fully three-dimensional flow description, and the more expensive turbulence closure required to represent anisotropic swirling-flow effects.

The refined 3D simulations require a separate interpretation. The baseline 3D solution was interpolated onto the refined mesh and used as the initial condition for the refined-mesh iterations; the effective cost of one refined-3D data point is therefore the sum of the average baseline-3D cost and the cost of the corresponding

refined-mesh run:

$$C_{\mathrm{3D,ref}} = C_{\mathrm{3D}} + C_{\text{refined only}} = 247.67 + 286.27 = 533.94 \text{ CPU h} \tag{4.56}$$

Each refined-mesh iteration required 5.71×10^{-1} CPU h, approximately four times the baseline-3D iteration cost of 1.42×10^{-1} CPU h. The values in Table 4.17 therefore represent the effective cost of the refined-mesh procedure adopted here; a cold-start of the refined simulation from a uniform initial field would have entailed a substantially higher total cost.

The cost increase across the hierarchy spans more than five orders of magnitude, from 5.0×10^{-3} CPU h for the CRN to 533.94 CPU h for the refined RANS–3D EARSM. This large separation confirms the relevance of a multi-fidelity strategy: the CRN and RANS–2D levels enable broad parametric exploration, whereas the three-dimensional levels provide higher-fidelity information at a substantially higher computational expense.

Table 4.17 Average computational cost of the different fidelity levels using the Naik chemical mechanism. The refined 3D cost is defined as the sum of the baseline 3D cost and the cost of the subsequent refined-mesh iterations; see Equation (4.56).

Fidelity level	Mean cost [CPU h]	Mean cost/iteration [CPU h/iter]
CRN, Naik	5.00×10^{-3}	–
RANS–2D, Naik	13.85	9.70×10^{-3}
RANS–3D EARSM, Naik	247.67	1.42×10^{-1}
RANS–3D EARSM refined, Naik	533.94	5.71×10^{-1}

Parametric Sensitivities

Table 4.15 isolates the effect of inlet temperature from that of residence time. At fixed mass flow ($\dot{m}_{\mathrm{air}} \approx 130$ kg/h) and $\phi = 0.40$, RANS–2D Naik predicts an increase from 3.09 ppm at $T_{\mathrm{in}} = 313$ K to 5.67 ppm at $T_{\mathrm{in}} = 693$ K, a factor of approximately 1.8; at $\phi = 0.60$ the same temperature span yields a factor of approximately 2.0 (from 13.52 to 26.86 ppm). The residence-time effect over the surveyed mass-flow range is of comparable or larger magnitude: at $T_{\mathrm{in}} = 453$ K and $\phi = 0.60$, RANS–2D Naik decreases by a factor of approximately 2.8 (from 25.30 to 9.02 ppm) as $\dot{m}_{\mathrm{air}}$ rises from 80 to 230 kg/h. The two effects therefore compete on a similar

quantitative footing; at the nominal operating points of Table 4.14, they reinforce one another, since T_{in} is increased while $\dot{m}_{\mathrm{air}}$ is reduced to maintain constant u_0, so that the apparent thermal sensitivity in that table is the combined signature of two distinct physical mechanisms.

A further qualitative difference emerges between architectures in the mass-flow sweeps. RANS–2D exhibits the strongest dependence on $\dot{m}_{\mathrm{air}}$, while RANS–3D EARSM is markedly less sensitive: at $T_{\mathrm{in}} = 453$ K and $\phi = 0.60$, the three-dimensional model decreases by only 25% (from 9.44 to 7.09 ppm) over a mass-flow range over which RANS–2D decreases by 64%. The CRN response is intermediate, with reductions of 32% and 43% for the Naik and CRECK mechanisms, respectively.

The inlet-temperature sensitivity reveals the same architectural ordering. At $\phi = 0.60$ and $\dot{m}_{\mathrm{air}} \approx 130$ kg/h, the predicted NO_x grows by a factor of approximately 4.5 between $T_{\mathrm{in}} = 313$ and 693 K in both RANS–3D EARSM (from 5.28 to 23.73 ppm) and CRN Naik (from 8.07 to 36.03 ppm), whereas RANS–2D Naik yields only a factor of approximately 2.0 (from 13.52 to 26.86 ppm); the same ordering is recovered at $\phi = 0.40$. With the kinetic source term held fixed across all three architectures, the spread in both sensitivities can be traced to the predicted residence-time and temperature fields rather than to the chemistry. For design exploration this implies that RANS–2D should be interpreted as an upper-bound estimator of NO_x at low $\dot{m}_{\mathrm{air}}$ within the investigated operating range, and that RANS–3D EARSM and the CRN provide more reliable guidance for extrapolation across the mass-flow envelope.

Velocity Field Validation

Before comparing emission predictions across fidelity levels, the fluid-dynamic soundness of the RANS–2D and RANS–3D EARSM simulations is assessed through a comparison of radial profiles of normalized axial (v_x/u_0) and radial (v_r/u_0) velocity at the reference operating condition ($\dot{m}_{\mathrm{air}} = 130$ kg/h, $T_{\mathrm{air}} = 623$ K, $\phi = 0.60$). The comparison includes the experimental PIV data [112] and LES results extracted from Amerighi et al. [5], which serve as an independent high-fidelity numerical reference. Profiles are shown at five axial stations: $x/D = 0.147$, 0.294, 0.588, 0.882, and 1.176, in Figs. 4.32–4.36.

At the two upstream stations ($x/D = 0.147$ and 0.294, near the combustion chamber entrance), the RANS–2D and RANS–3D EARSM predictions are in close mutual agreement for both velocity components. The main jet peak, the inner recirculation zone boundary, and the outer radial extent of the swirling flow are reproduced consistently between the two formulations. The LES profiles likewise cluster with the two RANS solutions. All four numerical predictions, however, deviate from the PIV profiles in this region: the measured axial velocity exhibits a more filled centreline distribution and a broader radial spreading of the jet, features that are not reproduced by any of the numerical approaches at these stations. This systematic discrepancy, shared by the LES, cannot be attributed to the limitations of the RANS approximation alone. Rather, it points to a structural CFD-to-experiment gap that persists regardless of the turbulence modelling fidelity. Possible contributing factors include uncertainties in the inlet turbulence intensity and length scale and minor discrepancies between the nominal combustor geometry and the actual test article.

Moving downstream, the level of agreement improves progressively. At $x/D = 0.588$ the jet peak location and the sign of the radial velocity are correctly captured by both RANS formulations, and at $x/D = 0.882$ all four predictions converge toward the PIV trend for both components. At the furthest station ($x/D = 1.176$), the RANS–3D EARSM profile is nearly coincident with the PIV measurements for both velocity components, while RANS–2D follows closely with a slight radial shift of the peak.

These results establish two conclusions relevant to the multi-fidelity emission analysis. First, the RANS–2D and RANS–3D EARSM simulations are internally consistent across the full axial extent of the combustion chamber: wherever the two solutions differ, the three-dimensional formulation is systematically closer to the experimental data, as expected from the more complete representation of swirl anisotropy provided by the EARSM closure. Second, the residual CFD-to-experiment gap observed near the chamber entrance is not specific to the present formulations; it is present at equal or greater magnitude in the reference LES [5], confirming that it reflects inherent limitations in replicating the experimental boundary conditions numerically rather than deficiencies of the RANS approaches adopted here.

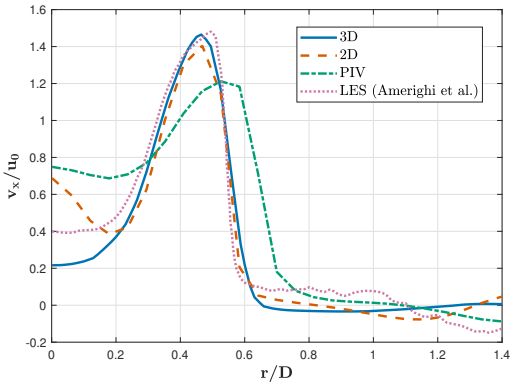

(a) Axial velocity

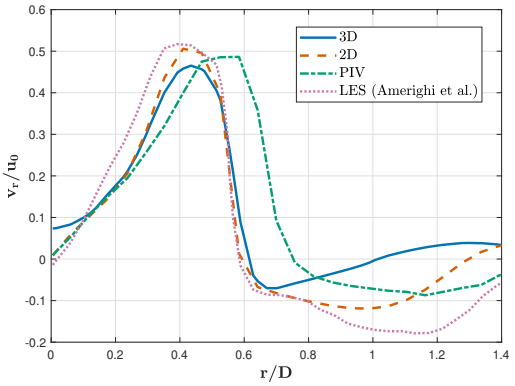

(b) Radial velocity

Fig. 4.32 Velocity profiles at $x/D = 0.147$.

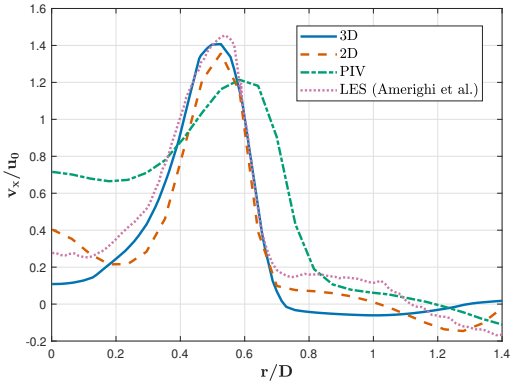

(a) Axial velocity

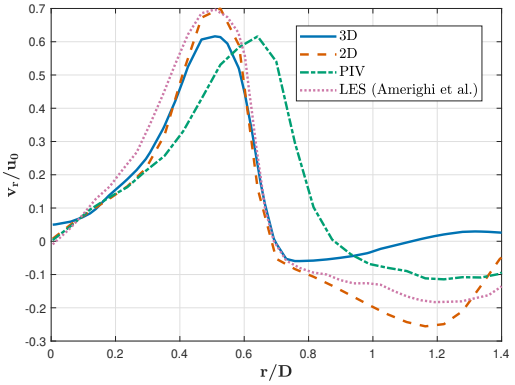

(b) Radial velocity

Fig. 4.33 Velocity profiles at $x/D = 0.294$.

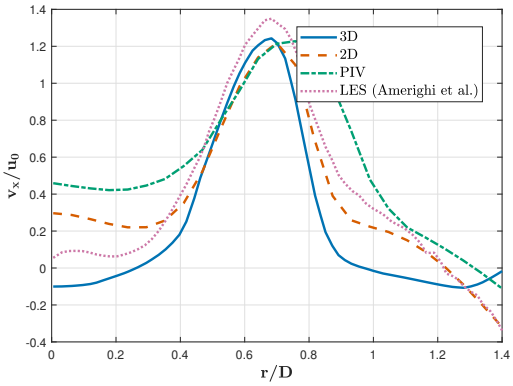

(a) Axial velocity

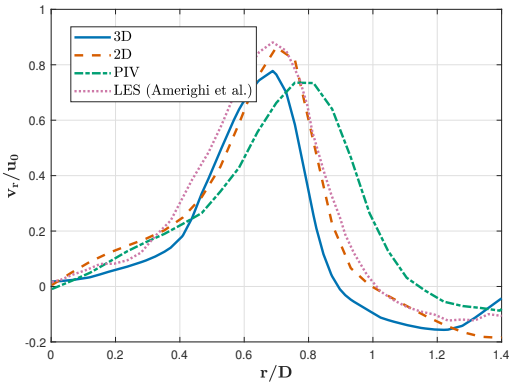

(b) Radial velocity

Fig. 4.34 Velocity profiles at $x/D = 0.588$.

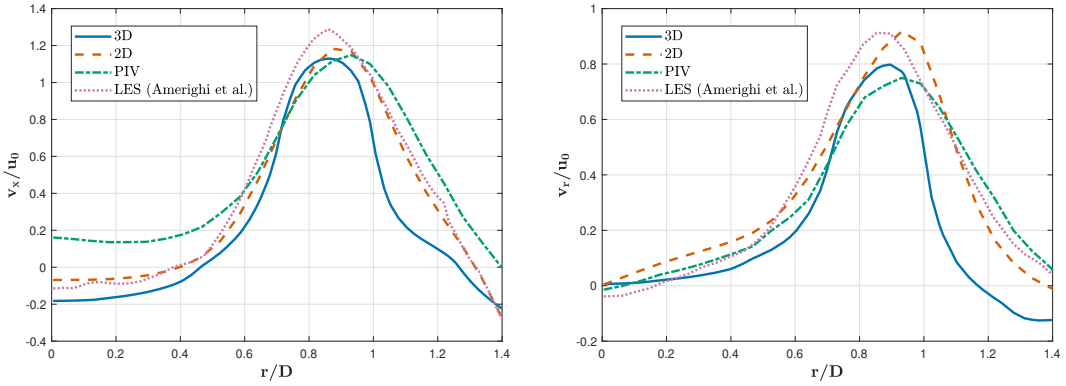

(a) Axial velocity (b) Radial velocity

Fig. 4.35 Velocity profiles at $x/D = 0.882$.

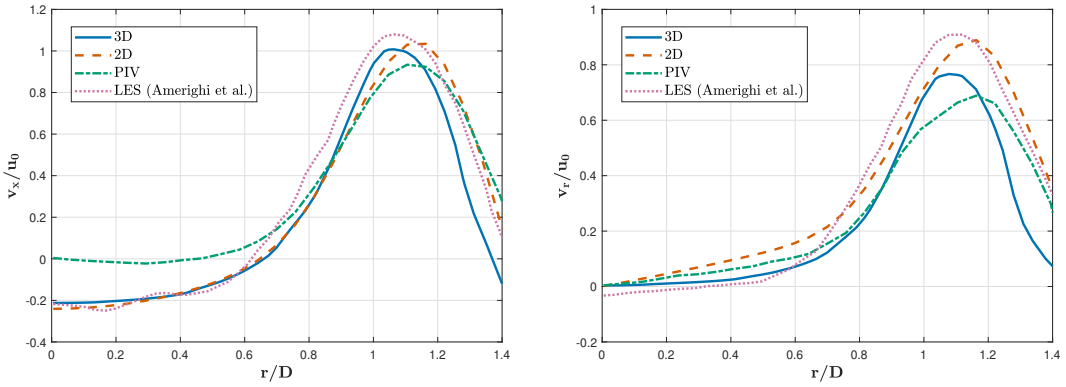

(a) Axial velocity (b) Radial velocity

Fig. 4.36 Velocity profiles at $x/D = 1.176$.

4.3.7 Uncertainty Quantification of NO_x Predictions

The calibrated CRN produces a deterministic NO_x prediction for each set of operating conditions, but several model inputs are known only within finite accuracy: the inlet flow rates and temperature carry measurement uncertainty, the calibrated mass-flow parameters retain residual ambiguity from the optimization process, and the geometric closure assumptions introduce further approximation. To characterise how these three sources of uncertainty propagate into the predicted NO_x output, and to identify which of them drives the predictive spread, a non-intrusive Polynomial Chaos Expansion (PCE) framework was applied to the calibrated network at the reference inlet condition [89].

The PCE acts as a fast surrogate of the CRN: a polynomial approximation of the NO_x response as a function of the uncertain inputs is first constructed from a moderate number of direct CRN evaluations, and is subsequently sampled at negligible computational cost. This strategy is necessary because a rigorous uncertainty analysis requires exploring the full space of possible input combinations: a task that would demand hundreds of thousands of direct CRN analyses, still expensive from a computational point of view. Once the surrogate has been built and validated, the same Monte Carlo sampling is performed on the polynomial approximation at virtually no cost, yielding the statistical moments and confidence intervals of the NO_x predictions.

An additional benefit of the PCE approach is that it provides, at no extra computational cost, variance-based Sobol sensitivity indices [135, 148]. These indices quantify the fractional contribution of each uncertain input to the total variance of the NO_x output: an index close to unity indicates that a given input alone accounts for nearly all of the output spread, while an index close to zero signals that the input has a negligible effect and can be safely excluded from further analysis. The distinction between first-order indices, which measure the direct effect of each input in isolation, and total-effect indices S_i^T, which additionally account for all interactions with other inputs, provides a complete picture of parameter relevance. When first-order and total-effect indices are nearly equal for all inputs, the model response is essentially additive and parameter interactions are negligible.

The quality of each surrogate is assessed through the Stone–Geisser leave-one-out predictive coefficient Q^2_{LOO} [145, 47, 18] . This metric measures how accurately the

surrogate predicts each training point when that point is temporarily withheld from the fitting procedure; values close to 1.0 are indicative of good predictive fidelity. A complementary assessment is provided by an independent external validation set, generated from a distinct sampling sequence and not used during fitting, through the coefficient of determination R^2_{test} and the root-mean-square error RMSE_{ext}.

Uncertainty Sources and Block Decomposition

The three classes of parametric uncertainty are organised into a block decomposition in which each class is first propagated independently through a dedicated low-dimensional PCE, and the results are then combined into a final joint block. Truncation to total degree p_o retains $P_{\text{basis}} = \binom{d+p_o}{p_o}$ basis functions; the configuration of the four blocks is summarised in Table 4.18.

Block O addresses the uncertainty in the inlet operating conditions of the AHEAD test rig [126]. The inlet temperature T_{air} is assigned a Gaussian distribution centred on the nominal value of 623 K with a standard deviation of 6.23 K, corresponding to an assumed 1% uncertainty on the temperature measurement chain. The relative deviations of the air and fuel mass-flow rates, δ_a and δ_f, are treated as independent zero-mean Gaussian variables with a standard deviation of 0.015, adopted as a conservative estimate of the flow-meter measurement uncertainty. The block therefore comprises three uncertain inputs ($d = 3$).

Block C characterises the local sensitivity of the NO_x output to small perturbations of the mass-flow calibration parameters θ_1–θ_8 around the reference vector $\vec{\Theta}^*$. Each of the eight parameters is modelled as a uniform variable over a narrow interval of half-width 0.025 centred on its optimised value. A preliminary correlation analysis on the ensemble of acceptable CMA-ES solutions ($J < 0.05$), using both Spearman rank and Pearson product-moment coefficients [139, 114], identified one strongly correlated pair, (θ_1, θ_8) [89]. To avoid violating the independence assumption underlying the PCE, θ_8 was fixed at its reference value, reducing the active set to seven parameters, as it had a lower Sobol index respect to θ_1. A variance-based screening step then applied the Sobol total-effect indices from an initial low-order surrogate to identify the five most influential parameters, whose cumulative contribution to the aggregated output variance exceeds 95%. Only these five are propagated to the final joint block.

Block G propagates the residual uncertainty associated with the geometric and closure assumptions of the CRN topology. The axial length L_{AX} of the Z_{AX} reactor is assigned a relative half-width of 20%, acknowledging the uncertainty associated with its definition. The Beér–Chigier swirl-jet spreading coefficient σ is assumed uniformly distributed over the interval $[0.60, 1.00]$ [15]. Volume fractions θ_9 and θ_{10} are treated as uniform variables with half-width 0.025 around their reference values, consistent with the perturbation treatment of Block C. The block therefore comprises four uncertain inputs ($d = 4$).

Block T combines the three operating inputs of Block O, the five retained calibration parameters from the Block C screening, and the four geometric/closure inputs of Block G into a single joint experimental design of twelve independent inputs ($d = 12$). This block provides the global predictive band on $\mathrm{NO_x}$.

All four blocks are propagated at the eight equivalence-ratio set points $\phi_{\mathrm{ref}} \in \{0.30, 0.40, 0.50, 0.60, 0.70, 0.80, 0.87, 0.96\}$ that span the experimental dataset of Reichel et al. [112]. At each set point, an experimental design is drawn from the input space using the Sobol low-discrepancy sequence [134] and propagated through the CRN. The polynomial coefficients are determined by least-squares regression with a sample size at least three times the number of basis functions, so that the system remains overdetermined even after any infeasible evaluations are discarded [18]. This is necessary as CRN simulations may fail to converge or may fail in enforcing the equivalence ratio control. A full Monte Carlo sampling of 2×10^4 realisations is then performed directly on the fitted surrogate to obtain the 95% uncertainty band from the empirical 2.5% and 97.5% quantiles.

Table 4.18 Configuration of the four PCE uncertainty blocks. d denotes the input dimensionality, p_o the polynomial order, P_{basis} the number of basis functions, N_{reg} the regression sample size per equivalence-ratio set point, and N_{ext} the size of the independent external validation set.

Block	d	p_o	P_{basis}	N_{reg}	N_{ext}
O (operating)	3	3	20	100	–
C (screening stage)	7	2	36	200	–
C (refinement stage)	5	3	56	500	150
G (volume/closure)	4	3	35	150	–
T (joint)	12	2	91	500	150

Results

The full uncertainty analysis required 14 000 direct CRN evaluations distributed across the four blocks and their external validation sets, executed in parallel over 48 CPU cores and completed in approximately 3.5 h.

Block O — Operating conditions. The surrogate mean reproduces the deterministic CRN curve within 0.5 ppm at every set point, confirming that small inlet perturbations are linear in the CRN output. The predicted standard deviation grows with the equivalence ratio, reaching a peak of $\sigma_O = 22.9$ ppm at $\phi = 0.87$ before decreasing slightly to 19.4 ppm at $\phi = 0.96$ (Fig. 4.39a). The mild decrease at the highest equivalence ratio is consistent with a saturation of the NO_x formation rate near-stoichiometric conditions, where the network response to small flow perturbations weakens.

The Sobol decomposition attributes the variance almost entirely to the two mass-flow perturbations: the ϕ-averaged total-effect indices are $\langle S_T \rangle_\phi = 0.537$ for δ_a and 0.420 for δ_f, whilst T_{air} contributes only 0.047. This result is physically expected, as both δ_a and δ_f act directly on the global equivalence ratio and thus on the adiabatic flame temperature. The inlet temperature sensitivity follows a U-shaped profile, peaking near the lean and stoichiometric extremes, but remains a secondary contributor throughout. First-order and total-effect Sobol indices coincide to within the third decimal for all three inputs, indicating that interactions among the operating uncertainties are negligible. Surrogate quality is high: $Q^2_{\text{LOO}} > 0.994$ at all eight set points.

Block C — Calibration parameters. The retention diagram in Fig. 4.37 summarises the result of the variance-based screening applied to the seven active calibration parameters. Two parameters dominate the output variance: θ_2, which controls the inlet air mass-flow partition between the two reservoirs ($\langle S_T \rangle_\phi = 0.362$, averaged over all the equivalence ratios), and θ_6, which governs the IRZ-to-FL_{SW} recirculation fraction (0.336). They are followed by θ_1 (0.153), θ_5 (0.101), and θ_4 (0.038). The remaining two parameters, θ_3 and θ_7, together contribute less than 4% of the aggregated score and are excluded from the joint analysis. This ranking is fully consistent with the CMA-ES calibration results of Section 4.3.4: the parameters that converge most tightly across seeds are precisely those with the highest Sobol contributions.

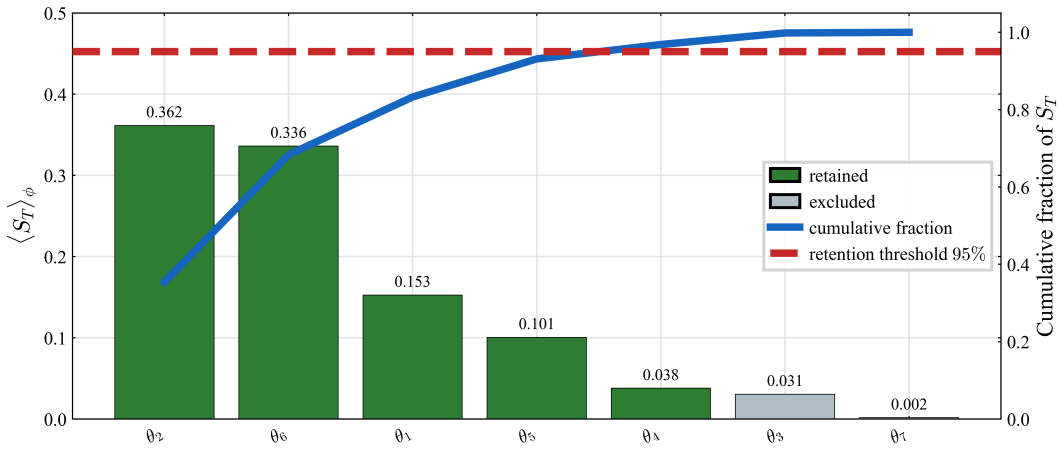

Fig. 4.37 Block C calibration retention diagram. Bars: ϕ-averaged total-effect Sobol score $\langle S_T \rangle_\phi$ from the screening surrogate. The cumulative fraction (right axis) crossing the 95% threshold (dashed line) defines the five retained parameters (green bars).

The refined Block C surrogate shows that the standard deviation induced by calibration uncertainty is one to two orders of magnitude smaller than that of Block O in the lean and intermediate regimes: σ_C grows from 0.008 ppm at $\phi = 0.30$ to 6.71 ppm at $\phi = 0.96$ (Fig. 4.39b). The leading parameter shifts with equivalence ratio: θ_6 controls the variance at lean and intermediate conditions, reflecting the dominant role of IRZ-to-swirl-zone recirculation in sustaining flame stability at low ϕ; above $\phi = 0.70$, the dominant contribution transitions to θ_2, whose total-effect index rises from 0.47 to 0.79 at $\phi = 0.96$. This transition may reflect a fluid dynamic change in the operating behaviour of the combustor. Surrogate accuracy is excellent, $Q^2_{\mathrm{LOO}} > 0.999$ at six of the eight set points, with lower but acceptable values at $\phi = 0.40$ (0.805) and $\phi = 0.50$ (0.808).

Block G — Geometric and closure inputs. The predicted standard deviation grows monotonically from $\sigma_G = 0.017$ ppm at $\phi = 0.30$ to 7.92 ppm at $\phi = 0.96$ (Fig. 4.39c), remaining below the operating contribution throughout the sweep and comparable in magnitude to the calibration block.

The ϕ-averaged Sobol attribution identifies the axial length L_{AX} as the dominant geometric uncertainty ($\langle S_T \rangle_\phi = 0.564$), since it directly sets the residence time in the Z_{AX} zone and thereby controls the extent of NO_x formation before the reacting mixture enters the swirl flame zone. The volume fraction θ_9 contributes 0.196, the swirl-jet spreading coefficient σ contributes 0.130, and θ_{10} contributes 0.118. The relative importance shifts with equivalence ratio: σ leads at $\phi = 0.40$ ($S_T = 0.45$)

where swirl-induced volume redistribution governs the early NO_x rise, whilst L_{AX} dominates from $\phi = 0.50$ onward. Near stoichiometric conditions, θ_9 and θ_{10} become comparable to L_{AX}, indicating that the downstream volume partitioning also contributes to the NO_x saturation. First-order and total-effect indices agree to within the third decimal throughout, confirming weak interactions among geometric inputs within the prescribed ranges.

Block T — Joint propagation and predictive band. The joint Block T surrogate is built on twelve independent inputs using an order-two polynomial expansion. External validation on $N_{ext} = 150$ fresh samples confirms its accuracy: $R^2_{test} \geq 0.999$ at seven of the eight set points, with $R^2_{test} = 0.994$ at $\phi = 0.40$. At the most demanding condition, $\phi = 0.96$, the external root-mean-square error remains below 0.42 ppm, corresponding to 0.14% of the nominal NO_x value. The resulting uncertainty statistics are reported in Table 4.19.

Table 4.19 Block T joint PCE statistics over the equivalence-ratio range. The mean μ_{NO_x}, standard deviation σ_{NO_x}, and 95% uncertainty band are obtained by Monte Carlo sampling of the surrogate. $RMSE_{ext}$ and R^2_{test} are evaluated on an independent validation set of $N_{ext} = 150$ samples per ϕ_{ref}.

ϕ_{ref}	μ_{NO_x} [ppm]	σ_{NO_x} [ppm]	CI^{lo}_{95} [ppm]	CI^{hi}_{95} [ppm]	$RMSE_{ext}$ [ppm]	R^2_{test} [–]
0.30	1.38	0.13	1.14	1.66	0.001	0.999
0.40	3.99	0.34	3.37	4.70	0.026	0.994
0.50	10.26	1.09	8.35	12.58	0.024	0.999
0.60	26.43	3.55	20.30	34.29	0.056	0.999
0.70	64.11	9.34	48.10	83.79	0.138	0.999
0.80	142.23	19.23	107.75	182.58	0.225	0.999
0.87	219.45	24.63	173.00	269.82	0.318	0.999
0.96	299.47	21.85	253.87	340.28	0.420	0.999

The aggregated Sobol decomposition of Block T, shown in Fig. 4.38, clarifies which uncertainty source drives the predictive band. The air and fuel mass-flow perturbations jointly account for approximately 87% of the total output variance ($\langle S_T \rangle_\phi = 0.49$ for δ_a and 0.38 for δ_f). The axial length L_{AX} and the inlet temperature T_{air} each contribute approximately 4%, whilst the volume fractions θ_9 and θ_{10} and the mass-flow partition θ_2 account for between 1% and 1.5% apiece. All remaining

calibration and geometric inputs lie below 0.5%. First-order and total-effect Sobol indices are very close for every input, confirming that the CRN NO_x response is essentially additive within the prescribed uncertainty ranges and that parameter interactions are negligible.

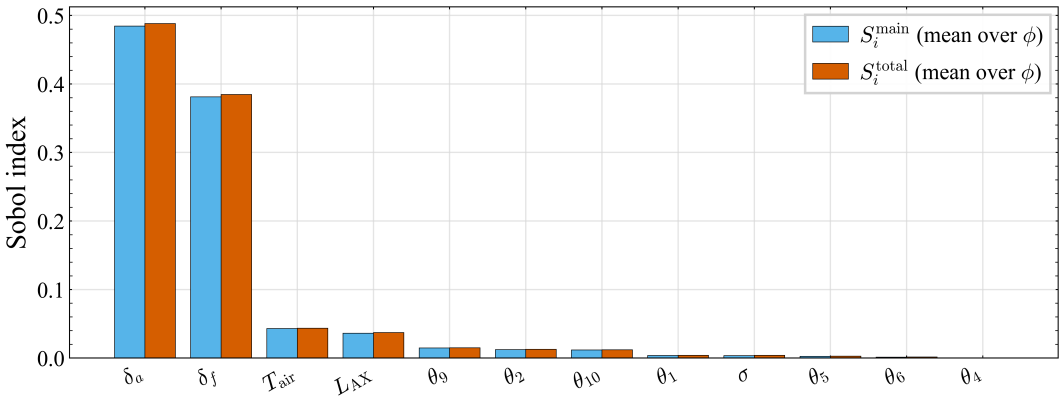

Fig. 4.38 Block T aggregated main (blue) and total-effect (orange) Sobol indices for the twelve independent inputs, averaged over the eight equivalence-ratio set points. The near-perfect overlap of the two bars for every input confirms that interactions among uncertainty sources are negligible.

The ϕ-dependence of the uncertainty apportionment follows a clear trend: the operating block accounts for 98.2% of the total-effect score at $\phi = 0.30$, declining to 90.3% at $\phi = 0.80$ and to 78.4% at $\phi = 0.96$. The geometric and closure inputs grow from 1.7% to 12.9% over the same range, driven primarily by L_{AX} and, near stoichiometric conditions, by the downstream volume fractions. The calibration parameters rise from below 1% in the lean regime to 9.3% at $\phi = 0.96$, governed by θ_2 and θ_1. The practical implication is clear: throughout the AHEAD operating envelope, the primary lever for reducing NO_x prediction uncertainty is the accuracy of the inlet flow-rate measurements, not the precision of the CRN calibration or geometric assumptions.

The block-by-block uncertainty bands and the final joint predictive band are presented in Figs. 4.39 and 4.40, respectively. The Block O envelope dominates at all equivalence ratios and the Block T band remains almost coincident with it, whilst the Block C and Block G bands are narrow throughout the lean and intermediate regimes.

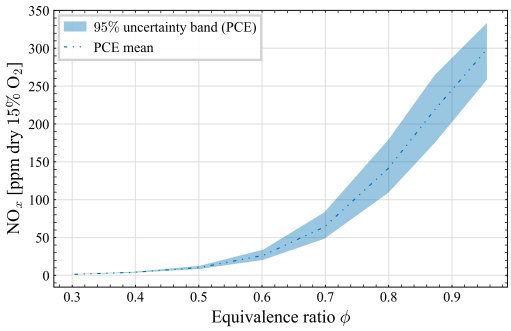

(a) Block O – operating conditions ($d = 3$).

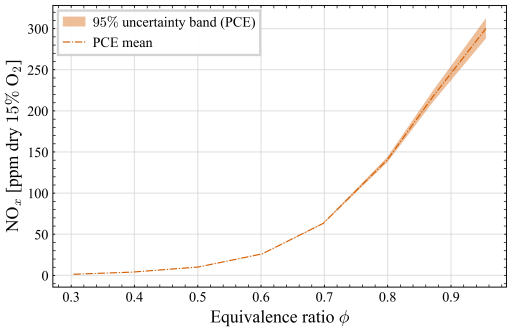

(b) Block C – retained calibration parameters ($d = 5$).

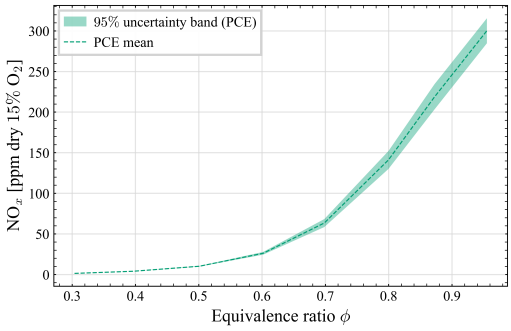

(c) Block G – geometric and closure inputs ($d = 4$).

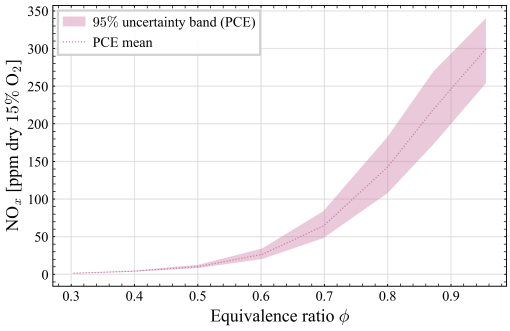

(d) Block T – joint propagation ($d = 12$).

Fig. 4.39 Block-by-block PCE mean (dashed line) and 95% uncertainty band (shaded region) on NO_x over the equivalence ratio range $\phi \in [0.30, 0.96]$. The operating block dominates the uncertainty throughout; calibration and geometric contributions become appreciable only near stoichiometric conditions.

The joint Block T band brackets six of the eight experimental measurements of Reichel et al. [126, 112], as shown in Fig. 4.40. The point at $\phi = 0.80$ lies at the upper edge of the predicted interval ($NO_{x,\mathrm{exp}} = 182.33$ ppm, band $= [107.75, 182.58]$ ppm). The two leanest set points, $\phi = 0.30$ and $\phi = 0.40$, fall above the upper bound of the band. This behaviour is consistent with the lean-limit limitations identified in Section 4.3.5: under very lean partially premixed hydrogen operation, localised hot spots and quenching–reignition events may generate NO_x levels that exceed those predicted by the homogeneous PSR network. These phenomena are not represented within the CRN model structure and therefore cannot be recovered by perturbing the model inputs within the prescribed uncertainty ranges.

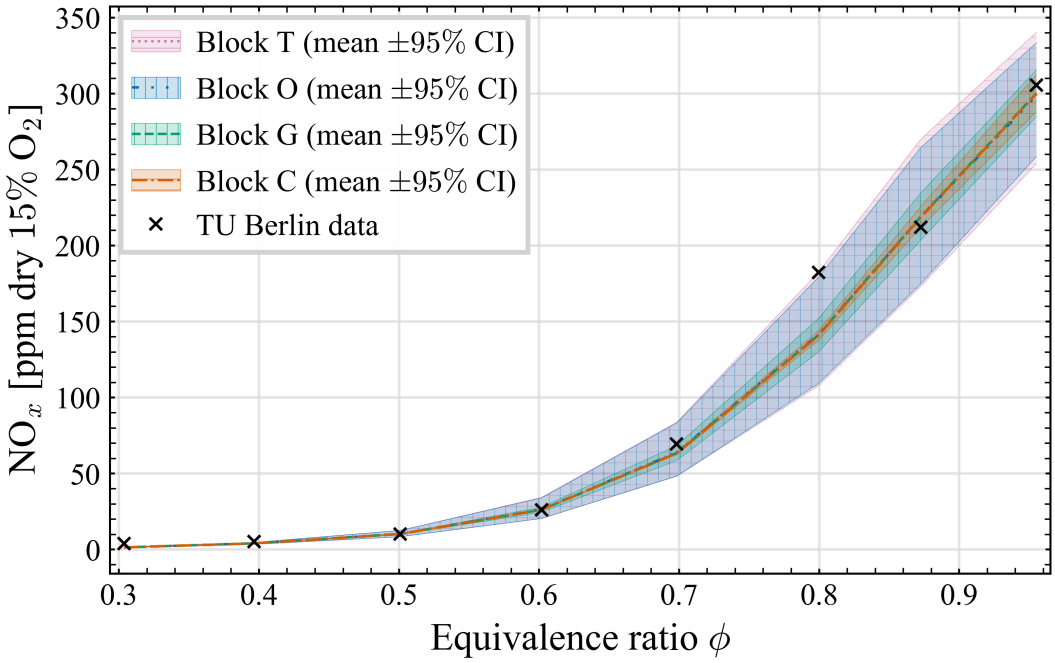

Fig. 4.40 95% NO_x uncertainty bands of the four UQ blocks overlaid on the same axes. The Block O and Block T bands are nearly coincident, confirming the dominance of inlet flow-rate uncertainty. The experimental data of Reichel et al. [126, 112] are shown as crosses. Six of the eight measurements are enclosed within the joint band; the two leanest points lie outside due to localised combustion phenomena inherently unresolved by the homogeneous network.

Overall, the calibrated CRN is robust against small perturbations of its optimised parameters and geometric assumptions. The residual predictive uncertainty on NO_x is dominated by inlet measurement accuracy throughout the lean and intermediate regimes, with calibration and geometric contributions becoming appreciable only near stoichiometric conditions. This result quantitatively supports the use of the CRN as a reliable and computationally efficient tool for hydrogen combustor emission analysis.

Summary of Case Study II

The AHEAD lean premixed combustor case study demonstrates that Chemical Reactor Network methodology can predict NO_x emission trends across wide ranges of operating conditions (inlet temperatures 313–693 K, mass flow rates 116–255 kg/h, equivalence ratios 0.30–0.96) with calibration performed at only a single operating

point. The distinguishing feature of this application is the swirl number-based volume allocation strategy, which reduces the need for preliminary CFD simulations while maintaining physical consistency with the characteristic flow structures.

The calibrated reactor network is validated against the full experimental matrix without recalibration, successfully reproducing the emission trends over the complete operating envelope within engineering accuracy bounds across both the lean and high-NO_x regimes.

The multi-fidelity comparative assessment contextualises the CRN within a three-level hierarchy that also includes two-dimensional axisymmetric RANS and fully three-dimensional RANS with Explicit Algebraic Reynolds Stress Model closure. The CRN achieves the lowest absolute and relative prediction error in the high-NO_x regime despite its substantially reduced model complexity, and remains competitive in the lean regime where all steady-state formulations are equally constrained by the absence of transient thermoacoustic phenomena. The added fidelity of the three-dimensional RANS formulation does not translate into consistently superior NO_x accuracy: its primary benefit lies instead in a more faithful reproduction of the three-dimensional velocity field and swirl anisotropy, as confirmed by the PIV velocity comparison. This fidelity–accuracy trade-off is accompanied by a cost separation exceeding five orders of magnitude across the hierarchy, from approximately 18 s per operating point for the CRN to approximately 534 CPU hours for the refined three-dimensional RANS simulation. This large separation confirms the value of a multi-fidelity strategy: the CRN enables broad parametric exploration at marginal cost, while the three-dimensional levels provide targeted high-fidelity information where required.

The Polynomial Chaos Expansion uncertainty quantification framework, organised into four independent blocks covering operating conditions, calibration parameters, and geometric closure assumptions, reveals that residual predictive uncertainty is dominated throughout the operating envelope by the accuracy of the inlet flow-rate measurements. Sobol sensitivity indices confirm that the two mass-flow perturbations jointly account for the vast majority of the total output variance, while calibration and geometric contributions become appreciable only near stoichiometric conditions. The joint 95% predictive band brackets six of the eight available experimental measurements; the two leanest points lie outside the band, consistent with localised transient phenomena that are inherently unresolved by the homogeneous

reactor network. This result carries a clear practical implication: the primary lever for reducing NO_x prediction uncertainty is experimental rather than computational, i.e. improving the accuracy of flow-rate measurements yields greater uncertainty reduction than refining the CRN calibration or geometric assumptions.

Overall, the case study validates the combination of simplified geometric and physical models with data-driven optimization to construct computationally efficient surrogate models for preliminary design and optimization in hydrogen aviation combustion systems. The swirl number-based volume allocation strategy demonstrates the generalisability of the CRN methodology to geometrically complex combustors, while the multi-fidelity and uncertainty quantification analyses provide a rigorous quantitative framework for assessing the reliability and applicability limits of the approach.

Chapter 5

Conclusions

The work conducted in this thesis investigated hydrogen combustion for aviation applications along three complementary and mutually reinforcing fronts: system-level propulsion analysis through the NPSS software framework, flame and fluid dynamic characterization via CFD simulations, and the development of a Chemical Reactor Network methodology for NO_x prediction employing a limited number of high-fidelity calibration data. Section 5.1 first summarises the three modelling fronts developed in the thesis. Section 5.2 then presents the principal results and achievements, and Section 5.3 outlines the most naturally motivated directions for future research.

5.1 Summary of Research Contributions

Three modelling approaches were developed, validated, and applied throughout this thesis, each addressing a distinct but interconnected aspect of hydrogen combustion in jet engines.

The first front, described in Chapter 2, concerned system-level thermodynamic analysis. Using the Numerical Propulsion System Simulation (NPSS) software, a full turbofan model including the cryogenic hydrogen fuel line, from the liquid storage tank to the combustor inlet, was constructed and parametrically studied [40]. The model benchmarked the baseline hydrogen-fuelled turbofan against its kerosene counterpart and evaluated two alternative thermodynamic cycles: a regenerative heat-exchange cycle, and an expander cycle. Crucially, cryogenic fuel system components,

including the tank, pump and heat exchangers were integrated into the engine model so that their thermodynamic interaction with the main cycle could be assessed in a self-consistent manner.

The second front, described in Chapter 3, concerned computational fluid dynamics. RANS simulations, 3D for the micromix injector [38], and both 2D axisymmetric and 3D for the AHEAD lean premixed swirl-stabilized combustor, were employed to characterize hydrogen flame behaviour, assess flashback risk, and establish the role of geometry on emission formation. The 2D axisymmetric formulation was coupled with a simplified multi-fidelity correction framework, SimpleMF [41], enabling accurate flame position prediction from a number of experimental training points comparable to the dimensionality of the input parameter space. A targeted LES of the AHEAD configuration provided high-fidelity reference data that simultaneously validated the 2D RANS approximation and revealed the three-dimensional turbulent structures that the lower-fidelity model necessarily smooths over.

The third and central front, described in Chapter 4, concerned the development of a Chemical Reactor Network (CRN) methodology for rapid NO_x prediction. The approach decomposes the combustor flow field into physically motivated zones, e.g. inner and outer recirculation regions, and treats the inter-reactor mass flow fractions as free parameters to be optimized against a minimal set of high-fidelity reference data. The methodology was validated on two substantially different configurations: a non-premixed swirl burner with water injection [39], and the AHEAD lean premixed combustor. The broader applicability of the framework was further demonstrated in Appendix A, where a CRN topology adapted to hybrid rocket engine combustion physics [43] was shown to reproduce experimental combustion efficiency data with errors consistently below or at 5 % across a wide range of operating regimes.

5.2 Principal Results and Achievements

5.2.1 Chemical Reactor Networks: Accuracy at Low Computational Cost

The CRN methodology constitutes the core scientific contribution of this thesis. Its most immediately quantifiable achievement is computational: for the AHEAD com-

bustor, a single CRN evaluation requires approximately 5.0×10^{-3} CPU h, compared to 13.85 CPU h for a 2D axisymmetric RANS simulation, and 533.94 CPU h for the 3D RANS procedure adopted in this work (see Table 4.17 in Chapter 4). The cost increase across the fidelity hierarchy therefore spans more than five orders of magnitude, from the CRN to the highest-fidelity RANS level, and the CRN efficiency would be highlighted even more if this comparison were made against LES.

This dramatic efficiency gain would be of limited value if obtained at the expense of predictive fidelity. The results presented in Sections 4.2 and 4.3 demonstrate that this is not the case. In Case Study I (non-premixed hydrogen burner with water injection), the CRN calibrated at the baseline condition, zero water loading, correctly captured the exponential NO_x reduction of 70–90 % observed under water injection, reproducing emission trends far from the calibration point without recalibration [39]. In Case Study II (AHEAD lean premixed combustor), a reactor network calibrated against four equivalence ratio measurements at a single reference condition predicted NO_x concentrations across the full experimental matrix, spanning inlet temperatures from 313 to 693 K, mass flow rates from 116 to 255 kg/h, and equivalence ratios from 0.30 to 0.96. Importantly, the CRN reproduced the experimental NO_x trends, at times even outperforming the much more computationally expensive CFD simulations in predicting the exact experimental data. It should be noted that this accuracy is conditional on the availability of calibration data: the CRN does not replace high-fidelity simulations but rather compresses their information content into a surrogate that can be interrogated rapidly across the design space. The robustness of these predictions was further assessed through a Polynomial Chaos Expansion (PCE) uncertainty analysis (Section 4.3.7), which identified inlet measurement accuracy as the dominant contributor to NO_x prediction variance, rather than the calibrated mass-flow parameters or the geometric volume allocation assumptions. This finding confirms that the network's predictive spread is physically grounded: rigorous characterization of combustor inlet conditions yields proportionally greater predictive confidence than refinement of the network architecture itself.

Taken together, these results establish the CRN as a reliable and practically useful tool for hydrogen combustor design, capable of delivering accuracy comparable to, and in certain regimes superior to, RANS simulations at a cost orders of magnitude lower. The computational efficiency of the CRN can be pushed further still through two complementary acceleration strategies developed in companion work [89], not reproduced here because they are accessory to the central contribution of the

thesis. Skeletal mechanism reduction via the Directed Relation Graph with Error Propagation (DRGEP) yields compact kinetic schemes with reduced stiffness, while In Situ Adaptive Tabulation (ISAT) replaces repeated direct integrations with rapid table lookups whenever the local interpolation error remains within a prescribed tolerance. Applied jointly within the CMA-ES calibration loop, the two strategies reduce the median wall-clock time of a full calibration campaign by more than an order of magnitude relative to the detailed-mechanism, direct-integration baseline, with no statistically significant degradation in solution quality.

5.2.2 Multi-Fidelity Modelling: Making the Most of What Is Available

A recurring theme throughout this work is the need to extract maximum predictive value from scarce high-fidelity data. This is addressed in two complementary ways: at the CFD level through the SimpleMF multi-fidelity correction framework [41], and at the emissions level through the CRN calibration strategy described above.

The multi-fidelity comparison between 2D RANS, 3D RANS, and LES carried out for the AHEAD combustor yields an important result: the 2D axisymmetric RANS approximation, despite its geometric simplifications, captures the essential velocity field structure of the combustor with a fidelity that is qualitatively consistent with both the 3D RANS and the time-averaged LES. The central recirculation zone, the swirl-induced radial expansion, and the characteristic near-wall reverse flow are all reproduced correctly. Quantitative differences exist, particularly in the mixing tube region, where the 2D approximation cannot represent the inherently three-dimensional swirler geometry, but these differences are sufficiently small that the 2D model provides a meaningful low-cost surrogate for the full 3D flow. The validation of this claim against LES is an important modelling result in its own right: it demonstrates that well-constructed 2D axisymmetric RANS simulations, which can be run with a relatively low computational cost, retain sufficient physics to serve as the low-fidelity backbone of a multi-fidelity framework.

The SimpleMF framework then shows that combining this affordable low-fidelity model with as few as three to five experimental calibration points yields predictions competitive with or superior to all alternative surrogate modelling strategies tested, such as Co-Kriging and NARGP, at small training set sizes. The results presented here

thus serve as a case study for the utility of multi-fidelity approaches in combustion research, demonstrating that physically informed low-fidelity models provide a robust foundation for data-efficient design space exploration.

Moreover, the simultaneous availability of two-dimensional and three-dimensional solutions for the same combustor naturally raised the question of whether the systematic gap between them could be closed at negligible additional cost. A lightweight modified U-Net, a convolutional neural network architecture, was trained to learn a data-driven correction of the two-dimensional fields toward the corresponding planar fields of the three-dimensional simulation, as reported in co-authored work [95, 42]. This approach, lying at the intersection of multi-fidelity CFD and machine learning, opens a research direction that extends beyond the physics-based modelling framework of this thesis. The result nonetheless points to a general possibility: once higher-fidelity reference data are available, whether from a three-dimensional RANS, a time-averaged LES, or experimental PIV measurements, a two-dimensional axisymmetric model can serve as a correctable baseline that a trained network effortlessly lifts toward the higher-fidelity target.

5.2.3 Cross-Domain Applicability: Chemical Reactor Networks for Hybrid Rockets

Appendix A demonstrates that the CRN methodology, developed and validated in the context of hydrogen gas turbine combustors, transfers successfully to a fundamentally different propulsion environment: the hybrid rocket engine, where combustion occurs at the interface between a solid fuel grain and a gaseous oxidizer, under conditions of spatially distributed vaporization and markedly different mixing physics [43]. Adapting the methodology required a dedicated reactor topology reflecting the hybrid rocket's characteristic zones, a physically motivated mixing fraction model to account for vaporization-limited entrainment of paraffin droplets, and an optimization-based calibration against NASA Ames experimental data [70]. Despite these substantial differences from the gas turbine setting, the calibrated CRN reproduced experimental combustion efficiency values with relative errors below 5 % across fifteen test conditions.

This result carries a significance that extends beyond the hybrid rocket application itself. It demonstrates that the CRN framework developed in this thesis is not a

combustor-specific tool, but a general methodology that can be adapted, with physically motivated choices for reactor topology and volume allocation, to any propulsion device in which finite-rate chemistry, residence time effects, and incomplete mixing play a controlling role.

5.3 Future Research Directions

The work presented in this thesis opens several clearly motivated research directions, each representing a logical extension of the capabilities developed here. Two are singled out as particularly impactful.

5.3.1 Generalizing the CRN Methodology

The CRN methodology developed in Chapter 4 rests on a set of modelling assumptions whose validity was demonstrated within the operating envelopes of the two test cases examined. The most critical of these is the assumption that the flow topology, defined by the inner and outer recirculation zones and the flame anchoring region, remains qualitatively stable as operating conditions vary. Future work must deliberately challenge this assumption by applying the framework to configurations where the topology is documented to change, which would motivate the development of regime-aware architectures in which the network structure is selected dynamically on the basis of fluid dynamics criteria.

A second structural limitation is the assumption of perfectly stirred conditions within each reactor zone. This is explicitly identified in Section 4.3.5 as a source of error at ultra-lean equivalence ratios, where localised hot spots and quenching-reignition events produce NO_x levels that exceed those predicted by the homogeneous network, as confirmed by the PCE analysis in Section 4.3.7. Replacing the PSR assumption with a Partially Stirred Reactor (PaSR) model, in which the effective chemical source term is modulated by the ratio of the chemical to the turbulent mixing timescale, is a natural remedy. However, this requires a methodology for estimating the PaSR mixing timescale within each reactor zone without incurring a sharp increase in computational cost, as would be the case if LES simulations were required for this purpose.

Finally, the CRN was exploited in this work as a steady-state model, advanced to equilibrium by time-marching, but it could in principle be extended to represent thermoacoustic instabilities. A possible formulation is to impose small-amplitude harmonic perturbations on the calibrated steady-state network across a range of frequencies and extract the resulting heat release response, yielding an estimate of the flame transfer function. Whether this approach can capture the essential physics of thermoacoustic coupling in hydrogen combustors with sufficient fidelity remains an open question and merits dedicated investigation.

5.3.2 An Integrated CRN–NPSS Cycle Optimization Workflow

A valuable and promising future direction is the deep integration of the CRN methodology into the NPSS engine model, creating a unified workflow for simultaneous cycle optimization and emissions assessment. The combination of the two tools developed in this thesis naturally enables this integration, and its potential impact on the preliminary design process for hydrogen aviation propulsion is substantial. The proposed workflow proceeds as follows.

1. **Engine cycle definition.** The NPSS turbofan model developed in Chapter 2 defines the thermodynamic operating point of the combustor: inlet total pressure, total temperature, air mass flow rate, and fuel-to-air ratio as functions of thrust setting and flight condition.

2. **Combustor preliminary sizing.** Given the cycle operating point, a preliminary combustor geometry is derived using established correlations for hydrogen-fuelled gas turbine combustors, following methodologies such as those proposed in [99]. This step defines the combustor volume, air distribution between zones, and the swirl number, providing the geometric input required for CFD simulations.

3. **CRN model construction and insertion into NPSS.** A CRN for the preliminary combustor geometry is constructed following the methodology of Chapter 4, calibrated against a minimal set of CFD data at the design point, and inserted into the NPSS model as a reduced-order combustor module. Within the cycle model, the CRN replaces the simple temperature-rise combustor

element, allowing the engine cycle to be closed with explicit NO_x emission information at the combustor exit.

4. **Multi-objective cycle optimization.** With the CRN-augmented NPSS model in place, the engine cycle can be re-optimized over the full design space, including overall pressure ratio, bypass ratio, turbine inlet temperature, combustor pressure drop, and equivalence ratio distribution, not merely for minimum thrust specific fuel consumption (TSFC) but simultaneously for minimum NO_x emission index. The CRN's negligible computational cost per evaluation makes it ideally suited for embedding within a population-based or gradient-based multi-objective optimization loop, where the engine model may be queried tens of thousands of times during the search.

5. **Emissions-aware optimum selection.** The multi-objective optimization produces a Pareto front in the TSFC–NO_x space, from which a design optimum can be selected according to the applicable regulatory target or a broader environmental impact metric accounting for both direct and indirect effects. This step explicitly captures the trade-off between cycle efficiency and emissions, a trade-off that cannot be quantified by purely thermodynamic cycle analysis.

6. **High-fidelity CFD validation.** The optimized combustor geometry and operating point are submitted to a final validation campaign to verify the CRN-predicted emission index at the design condition and identify any flow features requiring geometric adjustment. If the CFD results differ substantially from the CRN predictions, the CRN is recalibrated against the new high-fidelity data and the optimization is re-run, completing an inner loop of refinement.

The value of this workflow lies in how it restructures the role of high-fidelity CFD within the design process. Under the conventional approach, CFD is used to evaluate a small number of candidate designs sequentially, with each evaluation consuming substantial computational resources and limiting the breadth of the design space that can be explored. Under the proposed workflow, CFD is used selectively for two purposes: to provide the training data that calibrate the CRN at the outset, and to provide a final validation of the selected optimum at the end. All intermediate design space exploration, which constitutes the computationally expensive part of the optimization, is handled by the CRN at negligible cost. In this light, the CRN methodology represents not merely an efficient emissions predictor, but a key

enabler for a fundamentally more capable approach to emissions-conscious hydrogen combustor design.

References

[1] David J Acheson. *Elementary Fluid Dynamics*. Oxford University Press, Oxford, 1990.

[2] Eytan J. Adler and Joaquim R. R. A. Martins. Hydrogen-powered aircraft: Fundamental concepts, key technologies, and environmental impacts. *Progress in Aerospace Sciences*, 141:100922, 2023. doi: 10.1016/j.paerosci.2023.100922. Special Issue on Green Aviation.

[3] Advanced Research Projects Agency - Energy. Hydrogen Steam and Inter-Cooled Turbine Engine (HYSITE), 2022. URL https://arpa-e.energy.gov/programs-and-initiatives/search-all-projects/hydrogen-steam-and-inter-cooled-turbine-engine-hysite. Accessed: March 2025.

[4] AHEAD Consortium. Advanced Hybrid Engines for Aircraft Development (AHEAD): Project Final Report. Collaborative Project Final Report ACP1-GA-2011-284636, European Commission, Seventh Framework Programme, March 2015. Project no. 284636, FP7, AAT.2011.6.1-2 Propulsion. Scientific Coordinator: Dr. Arvind Gangoli Rao, Delft University of Technology.

[5] M. Amerighi, A. Andreini, T. Reichel, T. Tanneberger, and C. O. Paschereit. LES investigation of a swirl stabilized technically premixed hydrogen flame with FGM and TFM models. *Applied Thermal Engineering*, 247:122944, 6 2024. doi: 10.1016/j.applthermaleng.2024.122944.

[6] *Ansys® Fluent, Release 2024R1 – Theory Guide*. ANSYS, Inc., Canonsburg, PA, USA, 2024.

[7] A. J. Aspden, M. S. Day, and J. B. Bell. Turbulence–flame interactions in lean premixed hydrogen: transition to the distributed burning regime. *Journal of Fluid Mechanics*, 680:287–320, 2011. doi: 10.1017/jfm.2011.164.

[8] Antonio Attili, Rachele Lamioni, Lukas Berger, Konstantin Kleinheinz, Pasquale E Lapenna, Heinz Pitsch, and Francesco Creta. The effect of pressure on the hydrodynamic stability limit of premixed flames. *Proceedings of the Combustion Institute*, 38(2):1973–1981, 2021. doi: 10.1016/j.proci.2020.06.091.

[9] Sergio Bagarello, Dario Campagna, and Ivano Benedetti. A survey on hydrogen tanks for sustainable aviation. *Green Energy and Intelligent Transportation*, page 100224, 2024. doi: 10.1016/j.geits.2024.100224.

[10] Ghobad Bagheri, Matteo Pelucchi, Alessandro Stagni, Alessio Frassoldati, Alberto Cuoci, Tiziano Faravelli, and Eliseo Ranzi. Comprehensive kinetic study of combustion technologies for low environmental impact: MILD and OXY-fuel combustion of methane. *Combustion and Flame*, 212:142–155, 2020. doi: 10.1016/j.combustflame.2019.10.026.

[11] Eva Balsa-Canto, Martin Peifer, Julio R. Banga, Jens Timmer, and Christian Fleck. Hybrid Optimization Method with General Switching Strategy for Parameter Estimation. *BMC Systems Biology*, 2:26, 2008. doi: 10.1186/1752-0509-2-26.

[12] A. Baroutaji, T. Wilberforce, M. Ramadan, and A. G. Olabi. Comprehensive Investigation on Hydrogen and Fuel Cell Technology in the Aviation and Aerospace Sectors. *Renewable and Sustainable Energy Reviews*, 106:31–40, 2019. doi: 10.1016/j.rser.2019.02.022.

[13] Timothy Barth and Dennis Jespersen. The design and application of upwind schemes on unstructured meshes. In *27th Aerospace Sciences Meeting*, Reno, NV, 1989. American Institute of Aeronautics and Astronautics. doi: 10.2514/6.1989-366.

[14] D. L. Baulch, C. J. Cobos, R. A. Cox, P. Frank, G. Hayman, Th. Just, J. A. Kerr, T. Murrells, M. J. Pilling, J. Troe, R. W. Walker, and J. Warnatz. Evaluated Kinetic Data for Combustion Modelling. Supplement I. *Journal of Physical and Chemical Reference Data*, 23(6):847–1033, 1994. doi: 10.1063/1.555953.

[15] János Miklós Beér and Norman A Chigier. *Combustion Aerodynamics*. Fuel and energy science series. Applied Science Publishers Limited, 1972.

[16] Ian H. Bell, Jorrit Wronski, Sylvain Quoilin, and Vincent Lemort. Pure and Pseudo-pure Fluid Thermophysical Property Evaluation and the Open-Source Thermophysical Property Library CoolProp. *Industrial & Engineering Chemistry Research*, 53(6):2498–2508, 2014. doi: 10.1021/ie4033999.

[17] Ali Cemal Benim and Khawar J. Syed. *Flashback mechanisms in lean premixed gas turbine combustion*. Academic Press, Waltham, Massachusetts, 1st edition, 2015. ISBN 0-12-800826-1.

[18] Géraud Blatman and Bruno Sudret. Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics*, 230 (6):2345–2367, 2011. ISSN 0021-9991. doi: 10.1016/j.jcp.2010.12.021.

[19] Joseph Bozzelli and Anthony Dean. O + NNH: A possible new route for NOX formation in flames. *Int. J. Chem. Kinet.*, 27:1097–109, 11 1995. doi: 10.1002/kin.550271107.

[20] SL Bragg. Application of Reaction Rate Theory to Combustion Chamber Analysis. Technical report, Combustion and Fuels Sub-Comittee of the Aeronautic Research Council, 1953.

[21] Joseph Brand, Sam Sampath, Frank Shum, Robert Bayt, and Jeffrey Cohen. Potential Use of Hydrogen In Air Propulsion. In *AIAA International Air and Space Symposium and Exposition: The Next 100 Years*, 11 2003. ISBN 978-1-62410-165-6. doi: 10.2514/6.2003-2879.

[22] G. D. Brewer, R. E. Morris, G. W. Davis, E. F. Versaw, G. R. Cunnington Jr, J. C. Riple, C.F. Baerst, and G. Garmong. Study of fuel systems for LH2-fueled subsonic transport aircraft. Volume 1. Final report, September 1976–December 1977. Technical report, Lockheed-California Co., Burbank (USA), 1978. URL https://ntrs.nasa.gov/api/citations/19780023142/downloads/19780023142.pdf.

[23] G. Daniel Brewer. *Hydrogen Aircraft Technology*. Routledge, New York, 1991. doi: 10.1201/9780203751480.

[24] Martin Buhmann. *Radial Basis Functions: Theory and Implementations*. Cambridge University Press, 2003. ISBN 9780521633383. doi: 10.1017/CBO9780511543241.

[25] Antonio Cavaliere and Mara de Joannon. Mild combustion. *Progress in Energy and Combustion Science*, 30(4):329–366, 2004. doi: 10.1016/j.pecs.2004.02.003.

[26] Cheon Hyeon Cho, Hee Sun Han, Chae Hoon Sohn, and Jeong Sik Han. Ignition Delay Time and Oxidation of a Kerosene Aviation Fuel and a Blended Jet50-Bio50 Fuel. *ACS Omega*, 6(40):26646–26658, October 2021. doi: 10.1021/acsomega.1c04002.

[27] European Commission, Directorate-General for Mobility, Transport, Directorate-General for Research, and Innovation. *Flightpath 2050 – Europe's vision for aviation – Maintaining global leadership and serving society's needs*. Publications Office, 2011. URL https://data.europa.eu/doi/10.2777/50266.

[28] Riccardo Concetti, Josef Hasslberger, Nilanjan Chakraborty, and Markus Klein. Analysis of water droplet interaction with turbulent premixed and spray flames using carrier phase direct numerical simulations. *Combustion Science and Technology*, 195(7):1411–1433, 2023. doi: 10.1080/00102202.2023.2182192.

[29] Riccardo Concetti, Josef Hasslberger, Nilanjan Chakraborty, and Markus Klein. Effects of liquid water addition on turbulent premixed hydrogen/air combustion. *Fuel*, 373:132314, 2024. doi: 10.1016/j.fuel.2024.132314.

[30] Riccardo Concetti, Josef Hasslberger, Thomas Sattelmayer, and Markus Klein. On the Chemical Effect of Steam Addition to Premixed Hydrogen Flames with Respect to NOx Emissions and Flame Speed. *Flow, Turbulence and Combustion*, pages 1–16 , 2024. doi: 10.1007/s10494-024-00551-5.

[31] Coordinating Research Council, Inc. Handbook of Aviation Fuel Properties. Technical Report CRC-530, Coordinating Research Council, Inc., Atlanta, GA, January 1982. URL https://apps.dtic.mil/sti/html/tr/ADA132106/index.html. Accession Number: ADA132106, Contract: DAAK70-81-C-0128.

[32] Nicholas Cumpsty. *Jet propulsion*. Cambridge University Press, 2003.

[33] F. Duchaine, T. Morel, and L. Y. M. Gicquel. Computational-Fluid-Dynamics-Based Kriging Optimization Tool for Aeronautical Combustion Chambers. *AIAA Journal*, 47(3):631–645, 2009. doi: 10.2514/1.37808.

[34] ENABLEH2 Consortium. Enabling Cryogenic Hydrogen-Based CO_2 Free Air Transport (ENABLEH2): Final Report. Research Project Final Report 769241, European Commission, Horizon 2020, November 2022. Grant Agreement No. 769241, Duration: September 2018–November 2022.

[35] Ivar S. Ertesvåg. Scrutinizing proposed extensions to the Eddy Dissipation Concept (EDC) at low turbulence Reynolds numbers and low Damköhler numbers. *Fuel*, 309:122032, 2022. doi: 10.1016/j.fuel.2021.122032.

[36] David B. Fenn, Loren W. Acker, and Joseph S. Algranti. Flight Operation of a Pump-Fed Liquid-Hydrogen Fuel System. Technical Memorandum NASA-TM-X-252, NASA Lewis Research Center, Cleveland, Ohio, April 1960.

[37] Marco Ferrarotti, Zhiyi Li, and Alessandro Parente. On the role of mixing models in the simulation of MILD combustion using finite-rate chemistry combustion models. *Proceedings of the Combustion Institute*, 37(4):4531–4538, 2019. doi: 10.1016/j.proci.2018.07.043.

[38] Lorenzo Folcarelli, Andrea Ferrero, Filippo Masseni, and Dario Pastrone. Numerical Investigation of Different Configurations for Hydrogen Fuelled Jet Engine Combustors. In *SciTech 2024 Forum*. American Institute of Aeronautics and Astronautics, 2024. doi: 10.2514/6.2024-0964.

[39] Lorenzo Folcarelli, Riccardo Concetti, Fabio Cozzi, Andrea Ferrero, Filippo Masseni, Dario Pastrone, Josef Hasslberger, and Markus Klein. A Chemical Reactor Network methodology for estimating NOx emissions under non-premixed lean hydrogen combustion with water injection. *International Journal of Hydrogen Energy*, 168:150561, 2025. doi: 10.1016/j.ijhydene.2025.150561.

[40] Lorenzo Folcarelli, Giacomo Gedda, Alessandra Zumbo, Andrea Ferrero, Filippo Masseni, and Dario Pastrone. Integrated Numerical Modeling of a

Hydrogen Turbofan Engine and Fuel System With Alternative Thermodynamic Cycles. In *SciTech 2025 Forum*. American Institute of Aeronautics and Astronautics, 2025. doi: 10.2514/6.2025-0328.

[41] Lorenzo Folcarelli, Anna Spagnolo, Fausto Dicech, Andrea Ferrero, Filippo Masseni, and Dario Pastrone. Multi-Fidelity Modeling of a Lean Premixed Swirl-Stabilized Hydrogen Burner With Axial Air Injection. In *SciTech 2025 Forum*. American Institute of Aeronautics and Astronautics, 2025. doi: 10.2514/6.2025-0941.

[42] Lorenzo Folcarelli, Sebastiano G. I. Marcì, Vincenzo Madonia, Andrea Ferrero, Filippo Masseni, and Dario G. Pastrone. Numerical Simulation of NOx Emissions in a Lean Premixed Hydrogen Combustor Using Different Levels of Fidelity. In *Aviation 2026 Forum*. AIAA, 2026. doi: 10.2514/6.2026-4794.

[43] Lorenzo Folcarelli, Filippo Masseni, and Dario G. Pastrone. Chemical Reactor Network for Hybrid Rocket Engines Optimization. In *SciTech 2026 Forum*. American Institute of Aeronautics and Astronautics, 2026. doi: 10.2514/6.2026-1574.

[44] Alexander Forrester, Andras Sobester, and Andy Keane. Multi-fidelity optimization via surrogate modelling. *Proc. R. Soc. A*, 463:3251–3269, 10 2007. doi: 10.1098/rspa.2007.1900.

[45] Harald H.-W. Funke, Nils Beckman, Jan Keinz, and Atsushi Horikawa. 30 Years of Dry Low NOx Micromix Combustor Research for Hydrogen-Rich Fuels: An Overview of Past and Present Activities. In *Volume 4B: Combustion, Fuels, and Emissions*, Turbo Expo: Power for Land, Sea, and Air, page V04BT04A069, 09 2020. doi: 10.1115/GT2020-16328.

[46] Guodong Gai, Sergey Kudriakov, Bernd Rogg, Abdellah Hadjadj, Etienne Studer, and Olivier Thomine. Numerical study on laminar flame velocity of hydrogen-air combustion under water spray effects. *International Journal of Hydrogen Energy*, 44(31):17015–17029, 2019. doi: 10.1016/j.ijhydene.2019.04.225.

[47] Seymour Geisser. The Predictive Sample Reuse Method with Applications. *Journal of the American Statistical Association*, 70(350):320–328, 1975. doi: 10.1080/01621459.1975.10479865.

[48] Sebastian Göke and Christian Oliver Paschereit. Influence of Steam Dilution on Nitrogen Oxide Formation in Premixed Methane/Hydrogen Flames. *Journal of Propulsion and Power*, 29(1):249–260, 2013. doi: 10.2514/1.B34577.

[49] Sebastian Göke, Sebastian Schimek, Steffen Terhaar, Thoralf Reichel, Katharina Göckeler, Oliver Krüger, Julia Fleck, Peter Griebel, and Christian Oliver Paschereit. Influence of Pressure and Steam Dilution on NOx and CO Emissions in a Premixed Natural Gas Flame. *Journal of Engineering for Gas Turbines and Power*, 136(9):091508, 2014. doi: 10.1115/1.4026942.

[50] David G. Goodwin, Harry K. Moffat, Ingmar Schoegl, Raymond L. Speth, and Bryan W. Weber. Cantera: An Object-oriented Software Toolkit for Chemical Kinetics, Thermodynamics, and Transport Processes. https://www.cantera.org, 2023. Version 3.0.0.

[51] Sanford Gordon and Bonnie J McBride. Computer Program for Calculation of Complex Chemical Equilibrium. *NASA Reference Publication (RP)*, 1994. URL https://ntrs.nasa.gov/api/citations/19950013764/downloads/19950013764.pdf.

[52] Gabriel J. Gotama, Akihiro Hayakawa, Ekenechukwu C. Okafor, Ryuhei Kanoshima, Masao Hayashi, Taku Kudo, and Hideaki Kobayashi. Measurement of the laminar burning velocity and kinetics study of the importance of the hydrogen recovery mechanism of ammonia/hydrogen/air premixed flames. *Combustion and Flame*, 236:111753, 2022. doi: 10.1016/j.combustflame.2021.111753.

[53] Inge R. Gran and Björn F. Magnussen. A Numerical Study of a Bluff-Body Stabilized Diffusion Flame. Part 2. Influence of Combustion Modeling And Finite-Rate Chemistry. *Combustion Science and Technology*, 119(1-6):191–217, 1996. doi: 10.1080/00102209608951999.

[54] L. L. Gratiet and J. Garnier. Recursive Co-Kriging Model for Design of Computer Experiments with Multiple Levels of Fidelity. *Int J Uncertain Quantif*, 4(5):365–386, 2014. doi: 10.1615/int.j.uncertaintyquantification.2014006914.

[55] A. K. Gupta, N. Syred, and D. G. Lilley. *Swirl Flows*. Energy and engineering science series. Abacus Press, Tunbridge Wells, Kent, 1984. ISBN 9780856261756, 0856261750.

[56] Alexander Görtz and Daniel Silberhorn. Thermodynamic Potential of Turbofan Engines with Direct Combustion of Hydrogen. In *Proceedings of the 33rd Congress of the International Council of the Aeronautical Sciences, Stockholm, Sweden*, pages 4–9, 2022.

[57] Nikolaus Hansen and Andreas Ostermeier. Completely Derandomized Self-Adaptation in Evolution Strategies. *Evolutionary Computation*, 9(2):159–195, 06 2001. doi: 10.1162/106365601750190398.

[58] Nikolaus Hansen, Sibylle D. Müller, and Petros Koumoutsakos. Reducing the Time Complexity of the Derandomized Evolution Strategy with Covariance Matrix Adaptation (CMA-ES). *Evolutionary Computation*, 11(1):1–18, 03 2003. doi: 10.1162/106365603321828970.

[59] R. K. Hanson and S. Salimian. Survey of Rate Constants in H/N/O Systems. In W. C. Gardiner Jr., editor, *Combustion Chemistry*, pages 361–421. Springer, New York, 1984. ISBN 978-0-387-90963-9.

[60] Josef Hasslberger, Gulcan Ozel-Erol, Nilanjan Chakraborty, Markus Klein, and Stewart Cant. Physical effects of water droplets interacting with turbulent premixed flames: a Direct Numerical Simulation analysis. *Combustion and Flame*, 229:111404, 2021. doi: 10.1016/j.combustflame.2021.111404.

[61] Josef Hasslberger, Riccardo Concetti, Nilanjan Chakraborty, and Markus Klein. Inertial effects on the interaction of water droplets with turbulent premixed flames: A Direct Numerical Simulation analysis. *Proceedings of the Combustion Institute*, 39(2):2575–2586, 2023. doi: 10.1016/j.proci.2022.07.039.

[62] G. Hordemann, M. Grätsch, T. Lingstädt, and T. Krieg. Aircraft with Fuel Cell Drive: Development and Testing of a Motor Glider. *ECS Transactions*, 35(32):3–12, 2011. doi: 10.1149/1.3645597.

[63] Ying Huang, Shih-Yang Hsieh, Hong-Gye Sung, and Vigor Yang. Large-Eddy Simulation of Combustion Dynamics of Lean-Premixed Swirl-Stabilized Combustor. *Journal of Propulsion and Power*, 19, 09 2003. doi: 10.2514/2.6194.

[64] Christian Igel, Nikolaus Hansen, and Stefan Roth. Covariance Matrix Adaptation for Multi-objective Optimization. *Evolutionary Computation*, 15(1):1–28, 03 2007. doi: 10.1162/evco.2007.15.1.1.

[65] International Atomic Energy Agency. Hydrogen as an Energy Carrier and Its Production by Nuclear Power. IAEA-TECDOC 1085, International Atomic Energy Agency, Vienna, May 1999. URL https://www.iaea.org/publications/5375/hydrogen-as-an-energy-carrier-and-its-production-by-nuclear-power.

[66] Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13(4):455–492, 1998. doi: 10.1023/A:1008306431147.

[67] T. Kadyk, C. Winnefeld, R. Hanke-Rauschenbach, and U. Krewer. Analysis and Design of Fuel Cell Systems for Aviation. *Energies*, 11(2):375, 2018. doi: 10.3390/en11020375.

[68] Sascha Kaiser, Oliver Schmitz, Paul Ziegler, and Hermann Klingels. The Water-Enhanced Turbofan as Enabler for Climate-Neutral Aviation. *Applied Sciences*, 12(23):12431, 2022. doi: 10.3390/app122312431.

[69] J. Kallo, G. Renouard-Vallet, G. Schmithals, J. Schirmer, and K. A. Friedrich. Fuel Cell System Development and Testing for Aircraft Applications – Lessons Learned from Fuel Cell APU, HY4, ENFICA-FC and HY-LITAS Projects. *ECS Transactions*, 86(13):39–47, 2018. doi: 10.1149/08613.0039ecst.

[70] M Arif Karabeyoglu, Greg Zilliac, Brian J Cantwell, Shane De Zilwa, and Paul Castelluci. Scale-up tests of high regression rate liquefying hybrid

rocket fuels. In *39th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, pages AIAA–2003–6475, 2003.

[71] Stanislav Karpuk, Yiyuan Ma, and Ali Elham. Design Investigation of Potential Long-Range Hydrogen Combustion Blended Wing Body Aircraft with Future Technologies. *Aerospace*, 10(6):566, 2023. doi: 10.3390/aerospace10060566.

[72] William M. Kays. Turbulent Prandtl Number—Where Are We? *Journal of Heat Transfer*, 116(2):284–295, 05 1994. doi: 10.1115/1.2911398.

[73] James Kennedy and Russell Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95 - International Conference on Neural Networks*, volume 4, pages 1942–1948. IEEE, 1995. doi: 10.1109/ICNN.1995.488968.

[74] M. C. Kennedy and A. O'Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13, 2000. doi: 10.1093/biomet/87.1.1.

[75] Alan Kéromnès, Wayne K. Metcalfe, K. Alexander Heufer, Neill Donohoe, Apurba K. Das, Chih-Jen Sung, and Henry J. Curran. An experimental and detailed chemical kinetic modeling study of hydrogen and syngas mixture oxidation at elevated pressures. *Combustion and Flame*, 160(6):995–1011, 2013. doi: 10.1016/j.combustflame.2013.01.001.

[76] Bhupendra Khandelwal, Yingchun Li, Priyadarshini Murthy, Vishal Sethi, and Riti Singh. Implication of Different Fuel Injector Configurations for Hydrogen Fuelled Micromix Combustors. In *Volume 4: Cycle Innovations; Fans and Blowers; Industrial and Cogeneration; Manufacturing Materials and Metallurgy; Marine; Oil and Gas Applications*, Turbo Expo: Power for Land, Sea, and Air, pages 293–298, 06 2011. doi: 10.1115/GT2011-46845.

[77] Diederik P. Kingma and Jimmy Lei Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*, 2015. doi: 10.48550/arXiv.1412.6980.

[78] G. W. Koroll and S. R. Mulpuru. The effect of dilution with steam on the burning velocity and structure of premixed hydrogen flames. In *Symposium (International) on Combustion*, volume 21, pages 1811–1819. Elsevier, 1988. doi: 10.1016/S0082-0784(88)80415-6.

[79] Michael Krejci, Olivier Mathieu, Andrew Vissotski, Sankaranarayanan Ravi, Travis Sikes, Eric Petersen, Alan Kéromnès, Wayne Metcalfe, and Henry Curran. Laminar flame speed and ignition delay time data for the kinetic modeling of hydrogen and syngas fuel blends. *Journal of Engineering for Gas Turbines and Power*, 135:021503, 02 2013. doi: 10.1115/1.4007737.

[80] Rachele Lamioni, Pasquale Eduardo Lapenna, Lukas Berger, Konstantin Kleinheinz, Antonio Attili, Heinz Pitsch, and Francesco Creta. Pressure-induced hydrodynamic instability in premixed methane-air slot flames. *Combustion Science and Technology*, 192(11):1998–2009, 2020. doi: 10.1080/00102202.2020.1768081.

[81] Chung K. Law. *Combustion Physics*. Cambridge University Press, 2006.

[82] Chung King Law, A. Makino, and T. F. Lu. On the off-stoichiometric peaking of adiabatic flame temperature. *Combustion and Flame*, 145:808–819, 2006. doi: 10.1016/j.combustflame.2006.01.009.

[83] David S. Lee, David W. Fahey, Piers M. Forster, Peter J. Newton, Ron C. N. Wit, Ling L. Lim, Bethan Owen, and Robert Sausen. Aviation and global climate change in the 21st century. *Atmospheric Environment*, 43:3520–3537, 2009. doi: 10.1016/j.atmosenv.2009.04.024.

[84] Zhiyi Li, Alberto Cuoci, and Alessandro Parente. Large Eddy Simulation of MILD combustion using finite rate chemistry: Effect of combustion sub-grid closure. *Proceedings of the Combustion Institute*, 37(4):4519–4529, 2019. doi: 10.1016/j.proci.2018.09.033.

[85] Cheng Lu, Linyao Zhang, Can Cao, Xiye Chen, Chang Xing, Hongqing Shi, Li Liu, and Penghua Qiu. The effects of N_2 and steam dilution on NO emission for a H_2/Air micromix flame. *International Journal of Hydrogen Energy*, 47(63):27266–27278, 2022. doi: 10.1016/j.ijhydene.2022.06.050.

[86] Cheng Lu, Linyao Zhang, Chang Xing, Li Liu, and Penghua Qiu. Effects of characteristic diameter, steam dilution, and equivalence ratio on NO formation for a H_2/Air micromix design. *International Journal of Hydrogen Energy*, 61:1133–1141, 2024. doi: 10.1016/j.ijhydene.2024.02.373.

[87] John K. Lytle. The Numerical Propulsion System Simulation: An Overview. NASA Technical Memorandum NASA/TM-2000-209915, NASA Glenn Research Center, Jun 2000. URL https://ntrs.nasa.gov/citations/20000063377.

[88] Yajin Lyu, A. J. Aspden, Linyao Zhang, Li Liu, and Penghua Qiu. Study of the chemical effect of steam dilution on NO formation in laminar premixed H_2/Air flame at normal and elevated pressure. *International Journal of Hydrogen Energy*, 46(24):13402–13412, 2021. doi: 10.1016/j.ijhydene.2021.01.133.

[89] Vincenzo Madonia, Lorenzo Folcarelli, Andrea Ferrero, Filippo Masseni, and Dario G. Pastrone. Uncertainty Quantification and Computational Cost Reduction of a Chemical Reactor Network for a Hydrogen Combustor. In *Aviation 2026 Forum*. AIAA, 2026. doi: 10.2514/6.2026-4632.

[90] Vincenzo Madonia, Lorenzo Folcarelli, Andrea Ferrero, Dario G. Pastrone, and Filippo Masseni. NOx Emission Prediction of a Lean Premixed Hydrogen Combustor via Reactor Network Analysis. In *SciTech 2026 Forum*. American Institute of Aeronautics and Astronautics, 2026. doi: 10.2514/6.2026-1186.

[91] Bjørn F. Magnussen. The eddy dissipation concept: A bridge between science and technology. In *ECCOMAS thematic conference on computational combustion*, volume 21, page 24. Lisbon, Portugal, 2005.

[92] Joseph Majdalani. Analytical Models for Hybrid Rockets. In Kenneth K. Kuo and Martin J. Chiaverini, editors, *Fundamentals of Hybrid Rocket Combustion and Propulsion*, chapter 5, pages 207–246. American Institute of Aeronautics and Astronautics, 2007. doi: 10.2514/5.9781600866876.0207.0246.

[93] P. C. Malte and D. T. Pratt. The Role of Energy-Releasing Kinetics in NOx Formation: Fuel-Lean, Jet-Stirred CO-Air Combustion. *Combustion Science and Technology*, 9(5-6):221–231, 1974. doi: 10.1080/00102207408960360.

[94] P. C. Malte and D. T. Pratt. Measurement of atomic oxygen and nitrogen oxides in jet-stirred combustion. *Symposium (International) on Combustion*, 15(1):1061–1070, 1975. doi: 10.1016/S0082-0784(75)80371-7. Fifteenth Symposium (International) on Combustion.

[95] S. G. I. Marcì, L. Muscarà, L. Folcarelli, A. Ferrero, F. Masseni, and D. G. Pastrone. Use of Convolutional Neural Networks for Multifidelity CFD Problems. In *Proceedings of the CEAS AIDAA Conference 2025*. Materials Research Forum LLC, 2026.

[96] Cecil Marek, Timothy Smith, and Krishna Kundu. Low Emission Hydrogen Combustors for Gas Turbines Using Lean Direct Injection. In *41st AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, Tucson, Arizona, July 2005. doi: 10.2514/6.2005-3776.

[97] Timothy Marquardt and Joseph Majdalani. Review of Classical Diffusion-Limited Regression Rate Models in Hybrid Rockets. *Aerospace*, 6(6), 2019. doi: 10.3390/aerospace6060075.

[98] T. R. Marrero and E. A. Mason. Gaseous Diffusion Coefficients. *Journal of Physical and Chemical Reference Data*, 1(1):3–118, 01 1972. doi: 10.1063/1.3253094.

[99] Jack D. Mattingly, William H. Heiser, and David T. Pratt. *Aircraft engine design*. AIAA education series, 2002.

[100] Bonnie J. McBride, Sanford Gordon, and Martin A. Reno. *Coefficients for calculating thermodynamic and transport properties of individual species*, volume 4513. NASA Langley Research Center, 1993. URL https://ntrs.nasa.gov/citations/19940013151.

[101] Cyrus B. Meher-Homji and Erik Prisell. Pioneering Turbojet Developments of Dr. Hans Von Ohain—From the HeS 1 to the HeS 011. *Journal of Engineering for Gas Turbines and Power*, 122(2):191–201, 2000. doi: 10.1115/1.483194.

[102] F. R. Menter. Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal*, 32(8):1598–1605, 1994. doi: 10.2514/3.12149.

[103] E. Merkl. Final report summary-ENOVAL (ENgine mOdule VALidators). *European Commission: Amsterdam, The Netherlands*, 2018. URL https://cordis.europa.eu/docs/results/604/604999/final1-enoval-final-reporting.pdf.

[104] Wayne K. Metcalfe, S. M. Burke, S. S. Ahmed, and Henry J. Curran. A hierarchical and comparative kinetic modeling study of C1–C2 hydrocarbon and oxygenated fuels. *International Journal of Chemical Kinetics*, 45(10): 638–675, 2013. doi: 10.1002/kin.20802.

[105] R. C. Mulready. Liquid Hydrogen Engines. In R. B. Scott, W. H. Denton, and C. M. Nicholls, editors, *Technology and Uses of Liquid Hydrogen*, chapter 5. Pergamon Press Ltd., New York, 1964.

[106] Chitralkumar V. Naik. Detailed mechanism for hydrogen combustion with NOx. Ansys, Inc., 2009. Originally developed by Reaction Design, now part of Ansys.

[107] NASA. Turbopump systems for liquid rocket engines. Special Publication (SP) NASA-SP-8107, NASA, 1974. URL https://ntrs.nasa.gov/api/citations/19750012398/downloads/19750012398.pdf.

[108] F. Nicoud and F. Ducros. Subgrid-Scale Stress Modelling Based on the Square of the Velocity Gradient Tensor. *Flow, Turbulence and Combustion*, 62(3): 183–200, 1999. doi: 10.1023/A:1009995426001.

[109] Carsten Olm, István Gy. Zsély, Róbert Pálvölgyi, Tamás Varga, Tibor Nagy, Henry J. Curran, and Tamás Turányi. Comparison of the performance of several recent hydrogen combustion mechanisms. *Combustion and Flame*, 161(9):2219–2234, 2014. doi: 10.1016/j.combustflame.2014.03.006.

[110] Optimad. romBOX, 2024. URL https://www.optimad.it/. Accessed: November 2025.

[111] Rahul Palulli, Kai Zhang, Simeon Dybe, Muhammad Yasir, Christian Oliver Paschereit, and Christophe Duwig. Investigation of NO formation in non-premixed, swirl-stabilised, wet hydrogen/air gas turbine combustor. *Fuel*, 378: 132943, 2024. doi: 10.1016/j.fuel.2024.132943.

[112] Christian Oliver Paschereit and Thoralf G. Reichel. AHEAD Lean-Premixed Swirl-Stabilized Hydrogen Burner TU Berlin Dataset. https://tnfworkshop.org/ahead-lean-premixed-swirl-stabilized-hydrogen-burnertu-berlin/, 2023. Accessed: May 2025.

[113] Suhas Patankar. *Numerical Heat Transfer and Fluid Flow*. CRC Press, Boca Raton, 1st edition, 1980. doi: 10.1201/9781482234213.

[114] Karl Pearson. VII. Note on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London*, 58(347-352):240–242, 12 1895. ISSN 0370-1662. doi: 10.1098/rspl.1895.0041.

[115] Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Review*, 60(3):550–591, 2018. doi: 10.1137/16M1082469.

[116] P. Perdikaris, M. Raissi, A. Damianou, N. D. Lawrence, and G. E. Karniadakis. Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2198):20160751, 2017. doi: 10.1098/rspa.2016.0751.

[117] Norbert Peters. *Turbulent Combustion*. Cambridge Monographs on Mechanics. Cambridge University Press, 2000.

[118] Bruce E. Poling, John M. Prausnitz, and John P. O'Connell. *The Properties of Gases and Liquids*. McGraw-Hill Professional, New York, 5th edition, 2001. ISBN 0-07-011682-2.

[119] J. W. Pratt and L. E. Klebanoff. Feasibility of the SF-BREEZE: A Zero-Emission, Hydrogen Fuel Cell, High-Speed Passenger Ferry. *International Journal of Hydrogen Energy*, 41(6):3255–3268, 2016. doi: 10.1016/j.ijhydene.2015.12.120.

[120] Anmol L. Purohit, Armen Nalbandyan, Philip C. Malte, and Igor V. Novosselov. NNH mechanism in low-NOx hydrogen combustion: Experimental and numerical analysis of formation pathways. *Fuel*, 292:120186, 2021. doi: 10.1016/j.fuel.2021.120186.

[121] Eliseo Ranzi, Alessio Frassoldati, Alessandro Stagni, Matteo Pelucchi, Alberto Cuoci, and Tiziano Faravelli. Reduced kinetic schemes of complex reaction systems: Fossil and biomass-derived transportation fuels. *International Journal of Chemical Kinetics*, 46(9):512–542, 2014. doi: 10.1002/kin.20867.

[122] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2006. ISBN 0-262-18253-X. doi: 10.7551/mitpress/3206.001.0001.

[123] Thoralf G. Reichel and Christian Oliver Paschereit. Interaction mechanisms of fuel momentum with flashback limits in lean-premixed combustion of hydrogen. *International Journal of Hydrogen Energy*, 42(7):4518–4529, 2017. doi: 10.1016/j.ijhydene.2016.11.018.

[124] Thoralf G. Reichel, Katharina Goeckeler, and Oliver Paschereit. Investigation of Lean Premixed Swirl-Stabilized Hydrogen Burner With Axial Air Injection Using OH-PLIF Imaging. *Journal of Engineering for Gas Turbines and Power*, 137(11):111513, 09 2015. doi: 10.1115/1.4031181.

[125] Thoralf G. Reichel, Steffen Terhaar, and Oliver Paschereit. Increasing Flashback Resistance in Lean Premixed Swirl-Stabilized Hydrogen Combustion by Axial Air Injection. *Journal of Engineering for Gas Turbines and Power*, 137 (7):071503, 07 2015. doi: 10.1115/1.4029119.

[126] Thoralf G. Reichel, Steffen Terhaar, and Christian Oliver Paschereit. Flashback Resistance and Fuel–Air Mixing in Lean Premixed Hydrogen Combustion. *Journal of Propulsion and Power*, 34(3):690–701, 2018. doi: 10.2514/1.B36646.

[127] Robert W. Ring, James E. Stott, and Christy Hales. Modeling the Risk of Fire/Explosion Due to Oxidizer/Fuel Leaks in the Ares I Interstage. In *3rd IAASS Conference*, Rome, Italy, October 2008. URL https://ntrs.nasa.gov/citations/20110004369. Report Number: M09-0064.

[128] G. Romeo, F. Borello, G. Correa, and E. Cestino. ENFICA-FC: Design of Transport Aircraft Powered by Fuel Cell & Flight Test of Zero Emission 2-Seater Aircraft Powered by Fuel Cells Fueled by Hydrogen. *International Journal of Hydrogen Energy*, 38(1):469–479, 2013. doi: 10.1016/j.ijhydene.2012.09.064.

[129] O. Schmitz, M. Hoorfar, and C. Minke. Fuel Cells and Hydrogen for Aircraft Auxiliary Power – Analysis, Status, Challenges and Solutions. *Journal of The Electrochemical Society*, 165(6):F580–F587, 2018. doi: 10.1149/2.1231807jes.

[130] H. J. Sheen, W. J. Chen, S. Y. Jeng, and T. L. Huang. Correlation of swirl number for a radial-type swirl generator. *Experimental Thermal and Fluid Science*, 12(4):444–451, 1996. doi: 10.1016/0894-1777(95)00135-2.

[131] Abe Silverstein and Eldon W. Hall. Liquid hydrogen as a jet fuel for high-altitude aircraft. Technical report, NACA-Lewis Flight Propulsion Laboratory, Cleveland, Ohio, 1955.

[132] Martin Skottene and Kjell Erik Rian. A study of NOx formation in hydrogen flames. *International Journal of Hydrogen Energy*, 32(15):3572–3585, 2007. doi: 10.1016/j.ijhydene.2007.02.038. International Symposium on Solar-Hydrogen-Fuel Cells 2005.

[133] Jasper Snoek, Hugo Larochelle, and Ryan Adams. Practical Bayesian Optimization of Machine Learning Algorithms. In F. Pereira, C.J. Burges, L. Bottou, and K. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25, pages 2951–2959. Curran Associates, Inc., 2012. URL https://proceedings.neurips.cc/paper_files/paper/2012/file/05311655a15b75fab86956663e1819cd-Paper.pdf.

[134] I.M Sobol'. On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics*, 7(4):86–112, 1967. ISSN 0041-5553. doi: 10.1016/0041-5553(67)90144-9.

[135] I.M Sobol'. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55 (1):271–280, 2001. ISSN 0378-4754. doi: 10.1016/S0378-4754(00)00270-6. The Second IMACS Seminar on Monte Carlo Methods.

[136] Y. Song, L. Marrodán, Nicolas Vin, Olivier Herbinet, Emmanuel Assaf, Christa Fittschen, Alessandro Stagni, Tiziano Faravelli, María U. Alzueta, and Frédérique Battin-Leclerc. The sensitizing effects of NO_2 and NO on methane low temperature oxidation in a jet stirred reactor. *Proceedings of the Combustion Institute*, 37(1):667–675, 2019. doi: 10.1016/j.proci.2018.06.115.

[137] V. Sosounov and V. Orlov. Experimental turbofan using liquid hydrogen and liquid natural gas as fuel. In *26th Joint Propulsion Conference*, Orlando, FL, USA, July 1990. American Institute of Aeronautics and Astronautics. doi: 10.2514/6.1990-2421.

[138] Southwest Research Institute. Numerical Propulsion System Simulation (NPSS), 2025. URL https://www.swri.org/markets/electronics-automation/software/aerospace-software/numerical-propulsion-system-simulation-npss. Managed by the NPSS Consortium. Version 3.2.0. Accessed: January 2026.

[139] C. Spearman. The Proof and Measurement of Association between Two Things. *The American Journal of Psychology*, 15(1):72–101, 1904. doi: 10.1037/11491-005.

[140] Alessandro Stagni, Alberto Cuoci, Alessio Frassoldati, Tiziano Faravelli, and Eliseo Ranzi. Lumping and reduction of detailed kinetic schemes: An effective coupling. *Industrial & Engineering Chemistry Research*, 53(22): 9004–9016, 2014. doi: 10.1021/ie403272f.

[141] Alessandro Stagni, Alessio Frassoldati, Alberto Cuoci, Tiziano Faravelli, and Eliseo Ranzi. Skeletal mechanism reduction through species-targeted sensitivity analysis. *Combustion and Flame*, 163:382–393, 2016. doi: 10.1016/j.combustflame.2015.10.013.

[142] Alessandro Stagni, Carlo Cavallotti, Suphaporn Arunthanayothin, Yu Song, Olivier Herbinet, Frédérique Battin-Leclerc, and Tiziano Faravelli. An experimental, theoretical and kinetic-modeling study of the gas-phase oxidation of ammonia. *Reaction Chemistry & Engineering*, 5(4):696–711, 2020. doi: 10.1039/C9RE00429G.

[143] Alessandro Stagni, Suphaporn Arunthanayothin, Mathilde Dehue, Olivier Herbinet, Frédérique Battin-Leclerc, Pierre Bréquigny, Christine Mounaïm-Rousselle, and Tiziano Faravelli. Low- and intermediate-temperature ammonia/hydrogen oxidation in a flow reactor: Experiments and a wide-range kinetic modeling. *Chemical Engineering Journal*, 471:144577, 2023. doi: 10.1016/j.cej.2023.144577.

[144] W. Steenbergen and J. Voskamp. The rate of decay of swirl in turbulent pipe flow. *Flow Measurement and Instrumentation*, 9(2):67–78, 1998. doi: 10.1016/S0955-5986(98)00016-8.

[145] M. Stone. Cross-Validatory Choice and Assessment of Statistical Predictions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2):111–133, 01 1974. ISSN 0035-9246. doi: 10.1111/j.2517-6161.1974.tb00994.x.

[146] Rainer Storn and Kenneth Price. Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. *Journal of Global Optimization*, 11(4):341–359, 1997. doi: 10.1023/A:1008202821328.

[147] Jochen Ströhle and Tore Myhrvold. An evaluation of detailed reaction mechanisms for hydrogen combustion under gas turbine conditions. *International Journal of Hydrogen Energy*, 32(1):125–135, 2007. doi: 10.1016/j.ijhydene.2006.04.005.

[148] Bruno Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability Engineering & System Safety*, 93(7):964–979, 2008. ISSN 0951-8320. doi: 10.1016/j.ress.2007.04.002. Bayesian Networks in Dependability.

[149] Saurav Tiwari, Michael J. Pekris, and John J. Doherty. A review of liquid hydrogen aircraft and propulsion technologies. *International Journal of Hydrogen Energy*, 57:1174–1196, 2024. doi: 10.1016/j.ijhydene.2023.12.263.

[150] David J. J. Toal. Applications of multi-fidelity multi-output Kriging to engineering design optimization. *Structural and Multidisciplinary Optimization*, 66(6):125, 2023. doi: 10.1007/s00158-023-03567-z.

[151] Ioan Cristian Trelea. The particle swarm optimization algorithm: Convergence analysis and parameter selection. *Inf. Process. Lett.*, 85(6):317–325, 2003. doi: 10.1016/S0020-0190(02)00447-7.

[152] S.R. Turns. *An Introduction to Combustion: Concepts and Applications*. McGraw-Hill Series in Mechanical Engineering. McGraw-Hill, Boston, 2000.

[153] Bram van Leer. Towards the Ultimate Conservative Difference Scheme. V. A Second-Order Sequel to Godunov's Method. *Journal of Computational Physics*, 32(1):101–136, 1979. doi: 10.1016/0021-9991(79)90145-1.

[154] Guenther von Elbe and Bernard Lewis. *Combustion, Flames and Explosions of Gases*. Elsevier Science & Technology, United States, 1987. ISBN 0124467512.

[155] Stefan Wallin and Arne V Johansson. An explicit algebraic reynolds stress model for incompressible and compressible turbulent flows. *Journal of Fluid Mechanics*, 403:89–132, 2000. doi: 10.1017/S0022112099007004.

[156] Stefan Wallin and Arne V. Johansson. Modelling streamline curvature effects in explicit algebraic reynolds stress turbulence models. *International Journal of Heat and Fluid Flow*, 23(5):721–730, 2002. ISSN 0142-727X. doi: 10.1016/S0142-727X(02)00168-6.

[157] Andreas Westenberger. Liquid hydrogen fuelled aircraft-system analysis. *CRYOPLANE, The European Commission, Brussels, Belgium, Report No. GRD1-1999-10014*, 2003. URL https://www.fzt.haw-hamburg.de/pers/Scholz/dglr/hh/text_2004_02_26_Cryoplane.pdf.

[158] Christopher Winnefeld, Thomas Kadyk, Boris Bensmann, Ulrike Krewer, and Richard Hanke-Rauschenbach. Modelling and Designing Cryogenic Hydrogen Tanks for Future Aircraft Applications. *Energies*, 11(1), 2018. doi: 10.3390/en11010105.

[159] Jiayang Xu, Cheng Huang, and Karthik Duraisamy. Reduced-Order Modeling Framework for Combustor Instabilities Using Truncated Domain Training. *AIAA Journal*, 58(2):618–632, 2020. doi: 10.2514/1.J057959.

[160] Puyan Xu, Changwei Ji, Shuofeng Wang, Xiaoyu Cong, Zedong Ma, Chuanqi Tang, Hao Meng, and Cheng Shi. Effects of direct water injection on engine performance in a hydrogen (H_2)-fueled engine at varied amounts of injected water and water injection timing. *International Journal of Hydrogen Energy*, 45(24):13523–13534, 2020. doi: 10.1016/j.ijhydene.2020.03.011.

[161] Y. B. Zeldovich. The Oxidation of Nitrogen in Combustion and Explosions. In J. P. Ostriker, G. I. Barenblatt, and R. A. Sunyaev, editors, *Selected Works of Yakov Borisovich Zeldovich, Volume I: Chemical Physics and Hydrodynamics*, chapter 25, pages 364–403. Princeton University Press, 1992. doi: 10.1515/9781400862979.364. Original article published in 1946.

[162] Y. B. Zeldovich. Oxidation of Nitrogen in Combustion and Explosions. In J. P. Ostriker, G. I. Barenblatt, and R. A. Sunyaev, editors, *Selected Works of Yakov Borisovich Zeldovich, Volume I: Chemical Physics and Hydrodynamics*, chapter 26, pages 404–410. Princeton University Press, 1992. doi: 10.1515/9781400862979.404.

[163] Victor P. Zhukov. Verification, Validation, and Testing of Kinetic Mechanisms of Hydrogen Combustion in Fluid-Dynamic Computations. *International Scholarly Research Notices*, 2012(1):475607, 2012. doi: 10.5402/2012/475607.

[164] Alessandra Zumbo, Andrea Ferrero, Filippo Masseni, and Dario Pastrone. Comparison of Methods for Thermodynamic Modelling of Cryogenic Storage Systems for Aerospace Applications. In *SciTech 2025 Forum*. American Institute of Aeronautics and Astronautics, 2025. doi: 10.2514/6.2025-0568.

[165] Marcus Ó Conaire, Henry J. Curran, John M. Simmie, William J. Pitz, and Charles K. Westbrook. A comprehensive modeling study of hydrogen oxidation. *International Journal of Chemical Kinetics*, 36(11):603–622, 2004. doi: 10.1002/kin.20036.

Appendix A

Application of Chemical Reactor Networks to Hybrid Rocket Engines

A.1 Introduction

Hybrid Rocket Engines (HREs) represent a propulsion technology significantly different from the hydrogen combustion systems discussed in the main body of this thesis. While hydrogen-fuelled gas turbines operate in continuous, premixed or non-premixed combustion regimes, hybrid rockets combine solid fuel with liquid or gaseous oxidizer, resulting in a diffusion-limited combustion process with inherently transient characteristics due to fuel grain regression. Despite these fundamental differences, both technologies share a common feature: they represent environmentally sustainable alternatives to conventional propulsion systems. Hybrid rockets using liquid oxygen and paraffin-based fuels produce primarily carbon dioxide and water vapour, avoiding the toxic emissions associated with hydrazine-based liquid propellants and the chloride compounds typical of solid rocket motors.

The purpose of this appendix is to demonstrate the versatility and broad applicability of the Chemical Reactor Network methodology developed in this thesis. The CRN approach, originally tailored for hydrogen combustion in gas turbines, can be successfully extended to other propulsion domains where finite-rate chemistry, mixing limitations, and residence time effects play critical roles in determining combustion efficiency. The application to hybrid rocket engines serves as a proof of concept that the same modelling philosophy can capture complex combustion phe-

nomena across different scales and configurations, from stationary power generation to aerospace propulsion.

A.2 CRN Topology for Hybrid Rocket Engines

The Chemical Reactor Network for hybrid rocket engines is constructed following the physical understanding of the flow field derived from classical hybrid rocket theory. The combustion chamber is discretised into four distinct Perfectly Stirred Reactors (PSRs), each representing a characteristic zone of the hybrid rocket flow field, as illustrated in Figure A.1.

The fuel-rich reactor represents the boundary layer region adjacent to the regressing paraffin surface. This zone contains both vaporized fuel and liquid droplets entrained from the molten paraffin layer through Kelvin-Helmholtz instabilities at the melt layer interface. The volume of this reactor is estimated as 15% of the turbulent boundary layer thickness over the grain length [97]. The flame reactor models the primary combustion zone located in the shear layer between the fuel-rich and oxidizer streams. The equivalence ratio in this reactor is maintained at a target value of 2 through feedback control of the incoming oxidizer flow, and its volume is estimated as 10% of the boundary layer volume [97]. The oxidizer-core reactor represents the oxidizer stream that has not yet mixed with fuel products, with its volume calculated as the remaining port cross-section after subtracting the fuel-rich and flame volumes. Finally, the mixer reactor models the post-combustion chamber where final mixing and burnout occur. The network topology is shown schematically in Figure A.2.

Four parameters define the mass flow connections between reactors and constitute the optimization vector. Parameter $\mathscr{A}$ represents the fraction of total oxidizer flow that diffuses into the fuel-rich boundary layer. Parameter $\mathscr{C}$ is the fraction of fuel that feeds the primary flame reactor. Parameter $\mathscr{G}$ controls the split of residual oxidizer between the mixer and direct bypass to the outlet, modelling the incomplete mixing typical of hybrid rockets where a portion of the oxidizer core exits through the nozzle without reacting. Parameter α is the droplet atomization factor that scales the liquid particle diameter, accounting for enhanced turbulent breakup beyond the critical Weber number criterion.

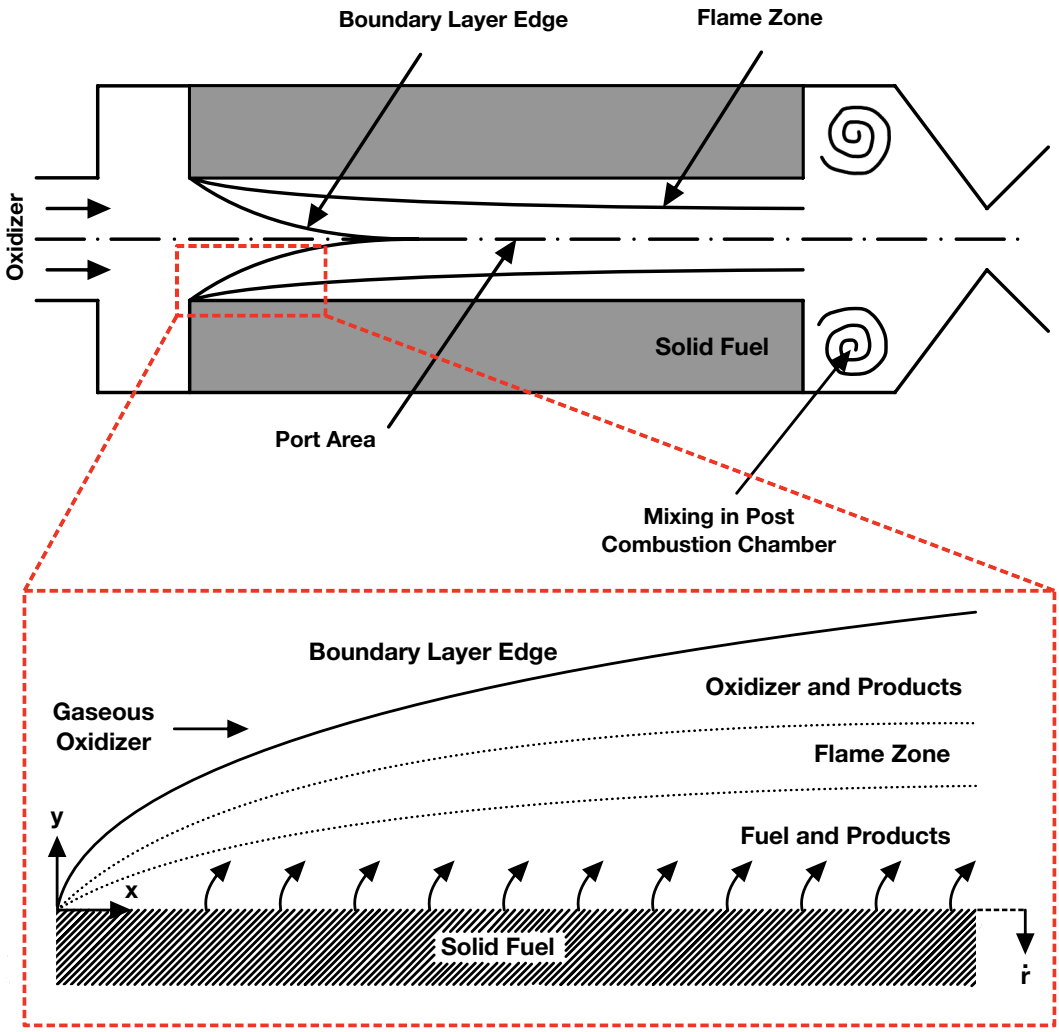

Fig. A.1 Schematic representation of a hybrid rocket engine combustion chamber. The upper portion shows the main components: oxidizer inlet, pre-chamber, fuel grain, port area, boundary layer, post chamber, and converging nozzle. The lower portion provides a detailed view of the fluid behaviour near the fuel grain, showing the oxidizer core, boundary layer edge, flame zone, fuel-rich zone, and the melting layer characteristic of paraffin fuels. Images adapted from Refs. [97, 92].

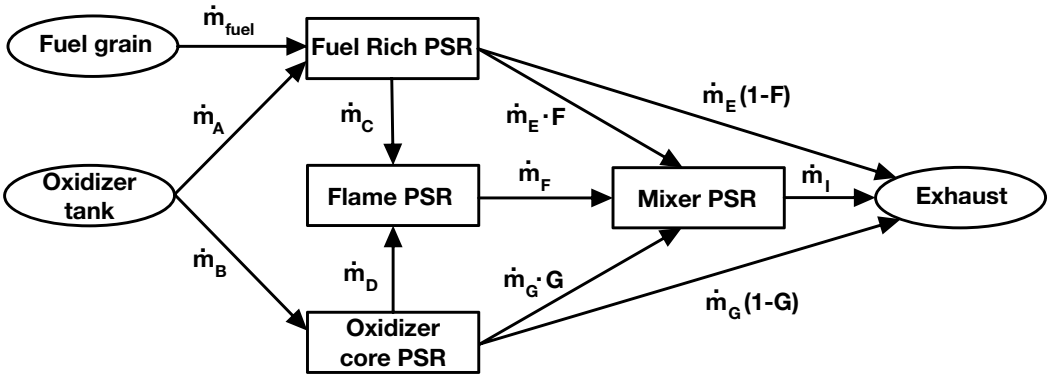

Fig. A.2 Chemical Reactor Network topology for hybrid rocket engine combustion chamber, showing the four perfectly stirred reactors and their interconnections through mass flow streams.

The oxidizer stream is initially split into diffusion and bulk components according to $\dot{m}_A = \mathscr{A} \cdot \dot{m}_{\text{ox,total}}$ and $\dot{m}_B = (1 - \mathscr{A}) \cdot \dot{m}_{\text{ox,total}}$. The fuel-rich reactor feeds the flame according to $\dot{m}_C = \mathscr{C} \cdot (\dot{m}_A + \dot{m}_{\text{fuel}})$, where the fuel mass flow rate is the product of regression rate, burning area, and paraffin density. The regression rate is computed using the empirical correlation $\dot{y} = a \cdot G_O^n$, and at each time step during the simulation the port area and all reactor volumes are updated to account for grain regression. The oxidizer flow to the flame is calculated from maintaining the target flame equivalence ratio, and the remaining flows are split according to the mixing fraction F and parameter $\mathscr{G}$ to model incomplete combustion and mixing.

Optimization is performed minimizing the error between predicted and experimental characteristic velocity through the fitness function

$$f = 1 - \sum_{i=1}^{N_{exp}} \left| \frac{\eta^*_{CRN,i} - \eta^*_{exp,i}}{\eta^*_{exp,i}} \right| \tag{A.1}$$

where the efficiency $\eta^*_{CRN,i}$ is computed by dividing the c^* obtained with the CRN by the ideal c^* obtained at equilibrium conditions with perfect mixing of all reactants. To facilitate CRN modelling and calculation of c_p using NASA polynomials, it is assumed that the oxygen is gaseous upon entering the combustion chamber. The ideal c^* is calculated using the same inital temperature for reactants to maintain consistency with the CRN model. The c^* inside the CRN is calculated considering the three outlet flows:

$$\Gamma_i = \sqrt{\gamma_i}\left(\frac{2}{\gamma_i+1}\right)^{\frac{\gamma_i+1}{2(\gamma_i-1)}} \tag{A.2}$$

$$K_{\mathrm{FR}} = \frac{\dot{m}_E\,(1-F)\sqrt{\dfrac{\mathscr{R}\,T_{\mathrm{FR}}}{\mathscr{M}_{\mathrm{FR}}}}}{\Gamma_{\mathrm{FR}}} \tag{A.3}$$

$$K_{\mathrm{Mixer}} = \frac{\dot{m}_I\sqrt{\dfrac{\mathscr{R}\,T_{\mathrm{Mixer}}}{\mathscr{M}_{\mathrm{Mixer}}}}}{\Gamma_{\mathrm{Mixer}}} \tag{A.4}$$

$$K_{\mathrm{OR}} = \frac{\dot{m}_G\,(1-G)\sqrt{\dfrac{\mathscr{R}\,T_{\mathrm{OR}}}{\mathscr{M}_{\mathrm{OR}}}}}{\Gamma_{\mathrm{OR}}} \tag{A.5}$$

$$p_t = \frac{K_{\mathrm{FR}} + K_{\mathrm{Mixer}} + K_{\mathrm{OR}}}{A_t} \tag{A.6}$$

$$c^* = \frac{p_t A_t}{\dot{m}_E\,(1-F) + \dot{m}_I + \dot{m}_G\,(1-G)} \tag{A.7}$$

where γ is the heat capacity ratio, A_t the throat area, $\mathscr{M}$ the molar mass, T the temperature and $\mathscr{R}$ the universal gas constant,

A.3 Mixing Fraction Model

A distinctive feature of the hybrid rocket CRN is the mixing fraction model, which addresses a physical phenomenon specific to paraffin-based hybrid propulsion. Unlike conventional hybrid fuels that rely solely on surface vaporization, paraffin combustion involves entrainment of liquid droplets from the unstable melt layer. These droplets must vaporize and burn before exiting the combustion chamber to achieve high combustion efficiency. Therefore, vaporization time becomes a rate-limiting process that must be explicitly modelled.

The mixing fraction F represents the fraction of the fuel-rich stream that effectively burns during the available residence time. The residence time in the combustion

chamber is calculated as $t_{\text{res}} = (V_{\text{total}}\rho)/\dot{m}_{\text{in}}$, where V_{total} is the sum of port volume and post-combustion chamber volume. The droplet diameter is estimated from critical Weber number scaling. Droplets become unstable and break up when the Weber number exceeds the critical value of approximately 12. The actual droplet diameter incorporates enhanced turbulent breakup through $d_{\text{droplet}} = \alpha\, d_{\text{Weber}}$, where the parameter $\alpha < 1$ quantifies the degree of atomization.

The vaporization time is calculated using the classical d^2-law as $t_{\text{vap}} = d^2_{\text{droplet}}/K_{\text{vap}}$, where the vaporization constant $K_{\text{vap}} = 8\lambda_{\text{gas}} \ln(1+B)/(\rho_{\text{F}} c_{p,\text{gas}})$ depends on gas thermal conductivity, fuel density, gas specific heat, and the Spalding transfer number $B = c_{p,\text{gas}}(T_{\text{eq}} - T_{\text{droplet}})/L_{\text{vap}}$. The expression used for B assumes negligible thermal inertia of the droplet [152]. Furthermore, for the subset of droplets burning, B should account for the heat of combustion per unit mass of fuel [152]: this effect is included for simplicity in the optimized value for α.

The fraction of fuel that actually enters the mixer reactor and burns is then calculated as

$$F = F_{\text{min}} + (1 - F_{\text{min}}) \frac{t_{\text{res}}/t_{\text{vap}}}{1 + t_{\text{res}}/t_{\text{vap}}} \tag{A.8}$$

where the parameter $F_{\text{min}} = 0.4$ represents a baseline combustion efficiency from direct vapour-phase fuel that does not undergo droplet entrainment. This formulation ensures correct asymptotic behaviour, with $F \to F_{\text{min}}$ as $t_{\text{res}}/t_{\text{vap}} \to 0$ and $F \to 1$ as $t_{\text{res}}/t_{\text{vap}} \to \infty$. This allows a portion of the fuel to burn even when the residence time is insufficient for complete droplet vaporization.

A.4 Experimental Data and Model Calibration

The CRN model for hybrid rockets was trained against experimental data from the NASA Ames Hybrid Combustion Facility test campaign, which investigated the high regression rate characteristics of paraffin-based fuels at operationally relevant conditions [70]. The test series employed fuel grains with 190.5 mm outside diameter and lengths of 838 or 1143 mm, operating at chamber pressures from 10.3 to 68.9 bar using liquid oxygen as the oxidizer. The fuel formulation SP-1a is a paraffin-based compound with a melting temperature of 342 K. The regression rate correlation obtained from the tests is $\dot{y} = 0.488\, G_O^{0.62}$ with $\dot{y}$ in mm/s and G_O in g/(cm^{2}s).

Five representative test conditions spanning the operational envelope were selected for CRN training. These tests covered chamber pressures from 18.3 to 44.3 bar, oxidizer mass flow rates from 2.05 to 4.45 kg/s, and both grain lengths available in the test series. The optimization procedure employed the Particle Swarm Optimization algorithm with a swarm of 10 particles evolved over 100 generations, minimizing the error between predicted and experimental characteristic velocity efficiency. The remaining experimental points were used to validate the CRN performance outside the training conditions.

The calibrated CRN model demonstrates excellent agreement with experimental data, achieving a mean absolute error of 1% on characteristic velocity efficiency across the five training points. The mean error on the validation points is approximately 1.5%. The physical properties of all training and validation test cases are summarized in Table A.1, while the comparison between experimental and predicted efficiencies is presented in Table A.2. The optimization procedure yields parameters $\mathscr{A} = 0.050$, $\mathscr{C} = 0.150$, $\mathscr{G} = 0.645$, and $\alpha = 0.080$. Computational efficiency is substantial, with CRN simulations requiring approximately 2×10^{-3} CPU-hours per case to simulate 0.1 s of engine operation.

Table A.1 Properties of experimental test cases from NASA Ames Hybrid Combustion Facility [70].

Test	P_c [bar]	$\dot{m}_O$ [kg/s]	O/F	$D_{\text{port, initial}}$ [m]	L_{grain} [m]	D_{throat} [m]
4F-1	36.3	2.03	3.05	0.1345	0.838	0.0371
4F-1a	34.5	4.02	3.89	0.1551	0.838	0.0503
4F-1b	38.7	2.13	2.72	0.1373	0.838	0.0368
4F-1c	37.4	2.07	3.00	0.1526	0.838	0.0368
4F-4	36.3	4.24	3.97	0.1345	0.838	0.0503
4F-5	38.0	4.32	3.54	0.1297	0.838	0.0503
4F-3a	39.2	4.39	3.84	0.1458	0.838	0.0503
4L-01	46.3	4.40	2.57	0.1440	1.143	0.0503
4L-03	44.3	4.45	2.69	0.1603	1.143	0.0503
4L-04	45.3	4.44	2.66	0.1239	1.143	0.0503
4L-05	44.7	4.43	2.72	0.1318	1.143	0.0503
4L-08	36.2	4.42	2.64	0.1341	1.143	0.0562
4L-09	18.3	2.05	1.70	0.1166	1.143	0.0556
4L-10	40.7	5.55	2.89	0.1480	1.143	0.0613
4L-12	20.8	2.08	2.01	0.1173	1.143	0.0513

Table A.2 Comparison of CRN model predictions with experimental data from NASA Ames Hybrid Combustion Facility [70]. Relative error defined as $\varepsilon = (\eta^*_{\text{CRN}} - \eta^*_{\text{exp}})/\eta^*_{\text{exp}} \cdot 100$.

Test	η^*_{exp}	η^*_{CRN}	ε [%]	Dataset
4F-1	0.85	0.864	+1.65	Training
4F-1a	0.91	0.892	−1.98	Validation
4F-1b	0.85	0.845	−0.59	Validation
4F-1c	0.89	0.845	−5.06	Validation
4F-4	0.90	0.899	−0.11	Validation
4F-5	0.88	0.880	+0.00	Validation
4F-3a	0.90	0.889	−1.22	Training
4L-01	0.88	0.846	−3.86	Validation
4L-03	0.84	0.849	+1.07	Validation
4L-04	0.85	0.854	+0.47	Training
4L-05	0.85	0.856	+0.71	Validation
4L-08	0.85	0.854	+0.47	Validation
4L-09	0.80	0.791	−1.13	Training
4L-10	0.88	0.863	−1.93	Validation
4L-12	0.79	0.794	+0.51	Training

A.5 Conclusions

This appendix demonstrates that the Chemical Reactor Network methodology developed for hydrogen combustion systems can be successfully adapted to hybrid rocket engines, despite the fundamentally different combustion regimes and physical phenomena involved. The key adaptations required include a reactor topology tailored to the characteristic flow zones of hybrid rockets, an optimization-based approach to determine mass flow distributions in the absence of detailed CFD data, and a physically-based mixing fraction model to account for the vaporization-limited combustion of entrained paraffin droplets.

The agreement between CRN predictions and experimental data, combined with the significant computational efficiency compared to full CFD simulations, validates the CRN approach as a versatile tool for combustion modelling across different propulsion applications.