

Hydrogen turbofan gradient-based optimisation including NO_x emissions constraints

Master of Science Thesis
Javier Alba Maestre



Hydrogen turbofan gradient-based optimisation including NO_x emissions constraints

Javier Alba Maestre

4844009



Supervisors: Prof. Joaquim R. R. A. Martins MDO Lab - University of Michigan
Dr. Andrew H. R. Lamkin MDO Lab - University of Michigan
Prof. Arvind Gangoli Rao Delft University of Technology

Institutions: Delft University of Technology Delft, The Netherlands
MDO Lab (University of Michigan) Ann Arbor, Michigan

Defence date: Friday 14th November, 2025

Cover Image: Boeing 787 Dreamliner at San Diego, Joaquim Martins, 2014



Summary

The debate around the the reduction of aviation emissions usually revolves around CO₂ emissions. The engine design trends of the last decades usually involve a an increase of the bypass ratio, overall pressure ratio, and turbine inlet temperature of engines. This trend has led to gains in efficiency and reductions in fuel consumption, consequently reducing carbon emissions. However, these developments have brought two main concerns. On the one side, these trends are hitting a plateau; limits in engine size or maximum allowable temperature of the turbine blades are reaching a barrier of how far this efficiency gain route can go. On the other hand, these developments result in higher combustor temperatures, which incur a penalty that can offset the benefits of CO₂ reduction: NO_x emissions. Increasing the combustor temperature beyond a certain point can lead to an exponential increase in NO_x emissions, which contribute to global warming, and pose health risk due to their toxicity. To solve this, new engine architectures or new fuels need to be considered, and NO_x emissions need to be part of the engine design cycle starting from its preliminary design.

To reach this goal, the research objective of this thesis is to

investigate the effects of including NO_x emission constraints in gradient-based optimisation of conventional and novel turbofan engines.

To achieve this, we decided to use simplified chemical reactor networks to predict emissions. Using simplified network architectures allows to use them during preliminary design without the need of detailed combustor or flow field geometry, posing a big advantage over CFD simulations. At the same time, these networks still allow for great flexibility in terms of inputs and operating conditions, so as to effectively analyse both conventional and unconventional designs with a level of accuracy that would not be possible with simpler approaches such as analytical correlations.

Two separate models are created in this work, one for conventional Jet A combustors, based on an architecture proposed by Villette et al. [1], and one for hydrogen lean premixed combustors, based on an architecture proposed by Talboom [2]. The former consists of three perfectly stirred reactors corresponding to the primary, secondary, and dilution zones of the combustor, and the later consists of four perfectly stirred reactors and one plug-flow reactor.

To use these models in a gradient-based optimisation framework, a surrogate model is created for each of them. The surrogate models are created using the Surrogate Modelling Toolbox [3], which provides analytical gradient information of the models. The models use the Kriging method, and the sampling dataset is created with Latin hypercube sampling.

To test the approach in an optimisation, a multipoint optimisation of a turbofan engine was carried out with pyCycle [4]. The baseline engine is based on the N+3 technology reference level by NASA [5], which consists of a geared, two-spool, high bypass turbofan. The engine was first optimised separately for both Jet A and hydrogen as fuels, to test the proposed methodology with conventional and unconventional fuels. Additionally, an engine with a closed-loop water recirculation system was optimised to test the approach with an unconventional architecture. For all of these cases, the optimisation was first performed without any emissions constraints, and then a second optimisation was carried out limiting the NO_x emissions to 60% of the baseline value for each case.

The multipoint optimisation problem consisted of one design point (top of climb), and three off-design points (take-off, sea level static, and cruise). The optimisation was performed using SNOPT [6], an SQP algorithm for large, constrained, gradient-based optimisation problems.

The takeaways of this research are the following:

- Simplified chemical reactor networks provide a good balance between accuracy and simplicity, allowing more freedom during preliminary design than empirical correlations. At the same time, they do not rely on detailed combustor geometry or flow field information that is necessary for more complex reactor network approaches. (See Chapter 3 and Chapter 4.)
- Surrogate modelling allows to integrate chemical reactor networks with gradient-based tools to overcome the lack of analytical gradients in tools like Cantera. They also offer other benefits, such as faster model probing, and avoid convergence issues with the reactor networks. (See Chapter 5.)
- The optimisations achieved 40% reduction in NO_x with a maximum penalty in fuel consumption of only 1.2%. This means that significant reductions in NO_x can be achieved with minimal penalties in fuel consumption, which supports the use of NO_x analysis and constraints during preliminary engine design. (See Chapter 7.)
- In the analysed engine configurations, H₂ leads to NO_x emissions that are 1 or 2 orders of magnitude below those of Jet A engines. (See Chapter 7.)
- Water recirculation can reduce NO_x emissions at the same time as improving engine performance. (See Chapter 7.)

Contents

Nomenclature	vi
List of Figures	ix
1 Introduction	1
1.1 Research Objective.	1
1.2 Research Questions	1
1.3 Project and Thesis Structure	2
2 Project Background	4
2.1 Climate Change and the Role of Aviation	4
2.2 Hydrogen in Aviation	6
2.2.1 Hydrogen Emissions versus Jet A Emissions.	6
2.2.2 Logistics	7
2.3 Novel Turbofan Engine Technologies	7
2.3.1 Challenges of Novel Architectures	8
2.3.2 Water Injection	9
2.4 Baseline Project	10
2.5 Framework and Tools	11
2.5.1 What is MDO?	11
2.5.2 OpenMDAO	12
2.5.3 Cycle Analysis: pyCycle	12
3 Emissions Modelling	13
3.1 Hydrogen Combustion	13
3.2 NO _x Emissions.	14
3.2.1 Thermal NO _x	15
3.2.2 Prompt NO _x	15
3.2.3 N ₂ O-intermediate NO _x	15
3.2.4 NNH Pathway	16
3.2.5 Fuel NO _x	16
3.2.6 NO _x Reburning	16
3.2.7 NO Oxidation to NO ₂	17
3.2.8 Effects of Steam on the Chemical Kinetics of NO _x Formation	17
3.3 Semi-Empirical Correlations	18

3.4	Chemical Reactor Networks	18
3.4.1	Cantera	19
3.4.2	Types of Reactors	19
3.5	Reaction Mechanisms	21
3.5.1	Luche's Mechanism for Jet A	21
3.5.2	Meng et al.'s Mechanism for H ₂ with NO _x pathways	21
4	Chemical Reactor Networks for Jet A and H₂ Combustors	23
4.1	Calculation of EINO _x	23
4.2	Jet A Combustor Model	23
4.2.1	Verification	25
4.2.2	Model Tuning	26
4.2.3	Extension to Water Recirculation	26
4.3	H ₂ Combustor Model	26
4.3.1	Verification	28
4.3.2	Extension to Full Size Combustor	31
5	pyCycle Integration through a Surrogate Model	33
5.1	Surrogate Modelling	33
5.1.1	Sampling: Latin Hypercube Sampling	34
5.1.2	Modelling: Kriging.	34
5.1.3	Surrogate Modelling Toolbox	35
5.2	Variable Scaling	35
5.3	Jet A Surrogate Models	35
5.3.1	Combustor Length Calculation	36
5.3.2	Cross-validation and Number of Sampling Points	36
5.4	H ₂ Surrogate Model	37
5.4.1	Cross-validation and Number of Sampling Points	38
5.5	Integration with pyCycle.	38
6	pyCycle Engine Model and Optimisation	39
6.1	Baseline Engine	39
6.2	Multipoint Problem Formulation	40
6.3	Optimisation Problem	41
6.3.1	Optimisation Cases	42
6.4	Optimiser	44

7	Optimisation Results and Comparison	45
7.1	Results of the Dry Cases.	45
7.2	Results of the Cases with Water Recirculation	46
7.3	Comparison of Cases.	48
7.4	Computational Time and Number of Iterations	49
8	Conclusion	52
9	Future Recommendations	54
	References	57
A	Sensitivity Analysis of a Cantera Reactor	58
A.1	Ignition of the Mixture and Effects of the Igniter Pulse Shape	58
A.2	Sensitivity of the Input Parameters	59
B	Hydrogen Engine with Water Recirculation	62
B.1	Extension of the Cantera Model for Water Recirculation	62
B.2	Surrogate Model.	62
B.3	Optimisation Results.	63

Nomenclature

Abbreviations

Abbreviation	Definition
BPR	Bypass Ratio
CFD	Computational Fluid Dynamics
CRN	Chemical Reactor Network
CRZ	Cruise
CS	Complex Step
DZ	Dilution Zone
EI	Emissions Index
FAR	Fuel-to-Air Ratio
FC	Flameless Combustion
FD	Finite Difference
GHG	Greenhouse Gas
HPC	High Pressure Compressor
HPT	High Pressure Turbine
LBO	Lean Blowout
LH2	Liquid Hydrogen
LHS	Latin Hypercube Sampling
LHV	Lower Heating Value
LNG	Liquefied Natural Gas
LPC	Low Pressure Compressor
LPT	Low Pressure Turbine
LTO	Landing and Take-off
MDO	Multidisciplinary Design Optimisation
MDP	Multi-Design Point
NPSS	Numerical Propulsion System Simulation
NO _x	Nitrogen Oxides
OPR	Overall Pressure Ratio
PFR	Plug-Flow Reactor
PSR	Perfectly Stirred Reactor
PZ	Primary Zone
RMSE	Root Mean Square Error
RTO	Rolling Take-Off
SLS	Sea-Level Static
SMT	Surrogate Modelling Toolbox
SNOPT	Sparse Nonlinear Optimizer
SQP	Sequential Quadratic Programming
SZ	Secondary Zone

Abbreviation	Definition
TIT	Turbine Inlet Temperature
TOC	Top Of Climb
TSEC	Thrust-Specific Energy Consumption
TSFC	Thrust-Specific Fuel Consumption
UHC	Unburned Hydrocarbons
UHR	Ultra High Bypass Ratio
XDSM	eXtended Design Structure Matrix

Symbols

Symbol	Definition	Unit
Latin symbols		
A	Area	m^2
AR	Air ratio	%
C_v	Nozzle velocity coefficient	—
c_p	Specific heat capacity at constant pressure	J/K/kg
D	Diameter	m
Da	Damköhler number	—
$dPqP$	Pressure loss	—
EI	Emissions index	g (of pollutant)/kg (of burnt fuel)
F	Thrust	N
FAR	Fuel-to-air ratio	—
H	Humidity factor	—
h	Specific humidity	kg (of water)/kg (of dry air)
h_i	Enthalpy of species i	J/kg
h^0	Enthalpy of formation	J/kg
J	Cost function	—
$K(x_i, x_j)$	Kernel function	—
L	Length	m
LHV	Lower heating value	J/kg
M	Mach number	—
\dot{m}	Mass flow	kg/s
\dot{m}'''	Mass production rate per volume	kg/s/m ³
MW	Molar mass	kg/kmol
p	Pressure	Pa
\dot{Q}	Heat transfer rate	W
R_u	Universal gas constant	J/kmol/K
T	Temperature	K

Symbol	Definition	Unit
t	Time	s
TSEC	Thrust-specific energy consumption	J/N/s
TSFC	Thrust-specific fuel consumption	kg/N/s
V	Volume	m^3
v	Velocity	m/s
v_x	Bulk flow velocity in axial direction	m/s
X_{H_2O}	Water fraction extracted from the exhaust flow	—
x	Axial position	m
Y_i	Mass fraction	kg (of species i)/kg (of total mixture)
$Z(x)$	Output of a stochastic process	*

Greek symbols

ζ	Heat loss fraction	—
η_c	Heat exchanger efficiency	—
θ	Surrogate model hyperparameter	—
μ_x	Mean value	Same as x
Π	Pressure ratio	—
ρ	Density	kg/m^3
σ_x	Standard deviation	Same as x
ϕ	Equivalence ratio	—
$\dot{\omega}$	Species production rate	$kmol/s/m^3$

Superscripts and Subscripts

$(\cdot)^*$	Active upper bound or \leq constraint
$(\cdot)^*$	Current value in the optimisation loop
$(\cdot)^{(0)}$	Initial guess for the optimiser
$(\cdot)_*$	Active lower bound or \geq constraint
$(\cdot)_3$	Last compressor exit/combustor inlet
$(\cdot)_4$	Combustor exit/turbine inlet
$(\cdot)_b$	Turbine blade
$(\cdot)_c$	Coolant
$(\cdot)_{ce}$	Coolant blade passage exit
$(\cdot)_{ci}$	Coolant blade passage inlet
$(\cdot)_{FL}$	Flight level
$(\cdot)_g$	Flue gas
$(\cdot)_i$	Species i
$(\cdot)_{SL}$	Sea level
$(\cdot)_x$	Axial direction

List of Figures

1.1	The thesis is divided into four main categories: emissions modelling, integration of the emission models with the engine optimisation tools, engine modelling, and optimisation.	2
2.1	Comparison of global warming forecasts depending on the preventive measures taken. Figure obtained from Ritchie, Rosado, and Roser [7].	5
2.2	Two most common configurations for aircraft with hydrogen propulsion. Image obtained from Adler and Martins [8].	5
2.3	Possible hydrogen tank configurations for tube+wing aircraft, where the darker shade of blue represents the passenger cabin. Image obtained from Adler and Martins [8].	5
2.4	Comparison of volumetric and mass energy densities of different energy carriers for aviation. The two fuels marked in red are the two analysed in this project. The data in the figure was obtained from Yin and Gangoli Rao [9].	6
2.5	Engine design trends for (a) bypass ratio, (b) overall pressure ratio, and (c) turbine inlet temperature.	8
2.6	Sample water injection diagram with water misting before the HPC and LPC. Image obtained from Daggett et al. [10].	9
2.7	Configuration of the two-spool turbofan with water recirculation.	11
3.1	Comparison of the flammability limits of hydrogen and kerosene. Figure obtained from Adler and Martins [8].	14
3.2	Diagram representing the different pathways for NO _x formation: thermal NO _x (green), prompt NO _x (red), N ₂ O-intermediate NO _x (blue), NNH-pathway NO _x (yellow), and fuel NO _x (purple). Image obtained from Iavarone et al. [11].	14
3.3	Representation of a perfectly stirred reactor (PSR). Image obtained from Turns [12].	19
3.4	Representation of a plug-flow reactor (PFR). Image obtained from Turns [12].	20
4.1	The chosen network for conventional Jet A combustors, based on the architecture proposed by Villette et al. [1], which consists of three reactors connected by valves, fed by reservoirs containing the oxidiser (air) and the fuel (Jet A), and features an igniter reservoir from which a pulse of hydrogen atoms flows to start the combustion reaction in the primary zone.	24
4.2	The reproduced model matches the temperature profile of Villette et al.'s architecture [1] within a maximum relative error of 0.6% across the three combustor zones.	25

4.3	The tuned data has air ratios and zone lengths common among the three engines, whereas Villette et al.'s data has different values for each engine. With the common parameters the NO _x emission trends are captured despite the error in the actual values being greater than the original for some of the cases. The data from Villette et al. was digitised from their figures [1].	27
4.4	The architecture proposed by Talboom [2] features 4 perfectly stirred reactors followed by a plug-flow reactor. There is recirculation between the flame, post-flame, and recirculation reactors. The reactors are fed by reservoirs containing the air and fuel, and ignition is modelled by a pulse of hydrogen atoms injected into the flame reactor.	27
4.5	Streamwise cross-section of the combustor. The vertical dimension corresponds to the diameter of the cylindrical combustor, and the horizontal direction is the length. The flow goes from left to right. The figure and the dimension definitions are based on the description from Talboom [2].	29
4.6	The results of the adiabatic verification cases show very good matching across different equivalence ratios, with the maximum error being within 0.3%. The simulated temperature is consistently higher across all points by 4/5 K.	29
4.7	The results of the verification cases with heat loss show good a good overall match, with a maximum relative error of 2.52% at the highest equivalence ratio. The absolute and relative errors are of the computed values (blue) with respect to the experimental data (black). The red values represent the results obtained by Talboom's CRN [2].	30
4.8	The comparison of the obtained NO _x trends show very good agreement with both the experimental (black) and reference (red) results from Talboom [2]. A noticeable underprediction of the NO _x levels occurs at the case with the lowest inlet temperature (313 K) combined with the highest equivalence ratio (0.9).	31
4.9	A plug-flow reactor is modelled as a series of perfectly stirred reactors, all having the same size. The dilution air is fed evenly to all PSRs. Instead of being connected to the rest of the network, the PFR region is modelled separately, where a reservoir with the properties of the post-flame region feeds the first PSR.	32
4.10	The full architecture features no heat loss to the environment, and a dilution stream connected to the PFR.	32
5.1	The RMS error of the Jet A model without water injection and its deviation starts to flatten at 21×n sampling points (126 total sampling points).	37
5.2	The RMS error of the Jet A model with water injection shows a minimum at 125 sampling points before increasing again.	37
5.3	The error of the H ₂ model without water injection starts to flatten at 24×n sampling points (120 total sampling points).	38
6.1	Architecture of the two-spool turbofan engine.	39
6.2	Architecture of the two-spool turbofan engine with the water recirculation system from Atma et al. [27].	40
6.3	Extended design structure matrix (XDSM) of the optimisation problem.	41

7.1	This figure shows the relative change in TSEC of each case with a NO _x constraint compared to its corresponding baseline case. Implementing NO _x constraints results in higher specific fuel and energy consumption except in case <i>JWNa</i> , where the addition of water during cruise offers an increase in efficiency while reducing emissions.	49
7.2	The addition of the surrogate model and the associated NO _x constraint results in more iterations for both dry cases.	50
7.3	The addition of the surrogate model and the associated NO _x constraint results in less iterations for the case in which the water extraction fraction is fixed (<i>JWNb</i>), but more iterations for the case in which the cruise extraction fraction is a design variable (<i>JWNa</i>).	50
A.1	The shape of the pulse affects the transient behaviour of the simulation, but does not affect the steady-state results as long as ignition is successfully achieved.	58
A.2	Increasing inlet temperatures result in higher exit temperatures, and might trigger autoignition in the reactor.	59
A.3	Higher pressures result in marginal gains in combustion temperature, but also hinder ignition due to the higher reactor mass.	60
A.4	Increasing the mass flow has a very small negative effect on exit temperature, and can also negatively affect ignition.	60
A.5	The equivalence ratio behaves as expected, with the highest temperature being around stoichiometric conditions.	61
A.6	The maximum temperature happens at an equivalence ratio slightly above 1, decreasing for both lower and higher equivalence ratios.	61
A.7	Increasing the combustor volume results in higher residence times and therefore more complete combustion, but increasing the reactor volume (and thus its mass content) beyond a certain point can make ignition unsuccessful.	61
B.1	The architecture of the CRN of the hydrogen combustor with water injection is the same as for the dry hydrogen combustor, with the addition of a water reservoir that feeds into the mixer reactor.	62
B.2	The error of the H ₂ model with water injection starts to flatten at 18 sampling points per variable (108 total sampling points).	63

Introduction

The debate around the the reduction of aviation emissions usually revolves around CO₂ emissions. The engine design trends of the last decades usually involve a an increase of the bypass ratio, overall pressure ratio, and turbine inlet temperature of engines. This trend has led to gains in efficiency and reductions in fuel consumption, consequently reducing carbon emissions. However, these developments have brought two main concerns. On the one side, these trends are hitting a plateau; limits in engine size or maximum allowable temperature of the turbine blades are putting a barrier into how far this efficiency gain route can go. On the other hand, these developments result in higher combustor temperatures, which carry a big penalty that can offset the CO₂ reduction: NO_x emissions. Increasing the combustor temperature beyond a certain point can lead to an exponential increase in NO_x emissions, which contribute to global warming, and pose health risk due to their toxicity.

To continue advancing aircraft propulsion towards a new era of performance and sustainability, new engine architectures can provide different was to improve efficiency and reduce emissions beyond what the current trends will allow. More over, engine design should account for NO_x emissions starting from the preliminary design stage, to ensure that new designs fall within emission reduction goals.

This project aims to provide a methodology to account for NO_x emissions for both conventional and novel architectures during the preliminary design stage.

1.1. Research Objective

The research objective of this thesis is to

investigate the effects of including NO_x emission constraints in gradient-based optimization of conventional and novel turbofan engines,

which will be achieved by

creating surrogate models with analytical gradients of chemical reactor networks of Jet A and hydrogen combustors, coupling them with pyCycle, and performing optimisations of Jet A and hydrogen turbofans, and a turbofan with water recirculation.

1.2. Research Questions

The research objective described in the previous section will be achieved by answering the following research questions:

1. How can emission prediction be integrated in a gradient-based optimisation framework for preliminary design of novel architectures without relying on empirical correlations?
 - (a) How can a chemical reactor network be constructed to model a generic gas turbine combustor during preliminary design without detailed geometry?
 - (b) How can the effects on emission trends of varying the design parameters be captured accurately with such a chemical reactor network?
 - (c) How can a chemical reactor network be coupled with pyCycle?
 - (d) How can the emission predictions from a chemical reactor network be used as objectives and constraints in an MDO problem of a gas turbine?

2. What are the effects in terms of thermodynamic performance, fuel consumption, and emissions of adding NO_x constraints to conventional and unconventional turbofan engines with conventional and unconventional fuels?
 - (a) How does the use of hydrogen fuel affect the optimised engine thermodynamic performance, fuel consumption, and emissions when compared to the baseline N+3 engine using Jet A?
 - (b) How does the addition of water recirculation to an N+3 turbofan affect the optimised engine thermodynamic performance, fuel consumption, and emissions when compared to the baseline N+3 engine?
 - (c) How does the addition of a NO_x constraint affect the three proposed architectures in terms of engine performance, emissions, and optimised engine configuration and dimensions?

1.3. Project and Thesis Structure

To help answer the proposed research questions, this thesis starts in Chapter 2 with a review of the emission goals that will constrain the future of aerospace propulsion. The chapter also contains relevant background information in terms of current design trends, novel technologies, and fuels, and outlines the baseline and framework that this project is built on.

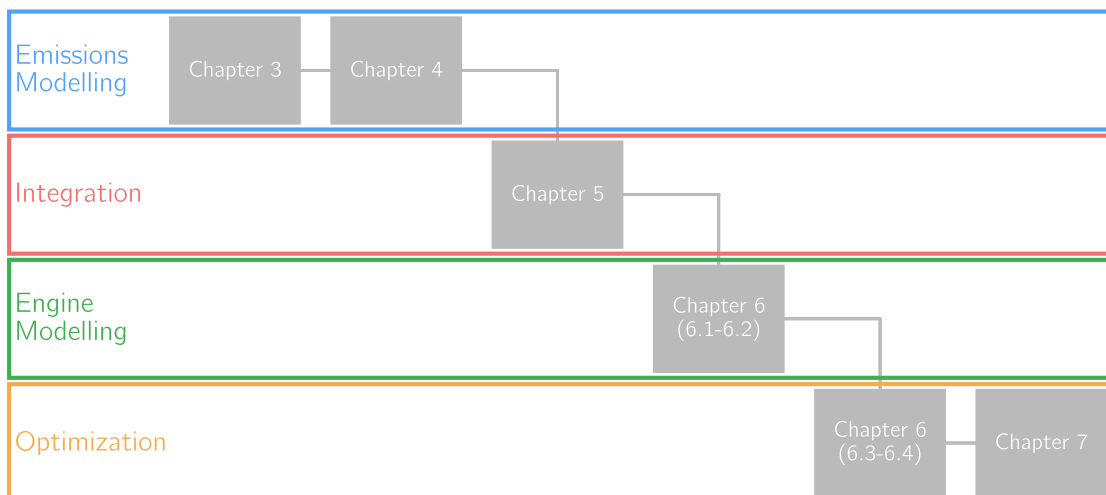


Figure 1.1: The thesis is divided into four main categories: emissions modelling, integration of the emission models with the engine optimisation tools, engine modelling, and optimisation.

Chapter 3 provides a comprehensive review of NO_x emissions, their formation mechanisms, and different approaches to predict them. Chapter 4 builds on this knowledge and describes the two models, one for Jet A and one for H₂, that will be used to predict emissions of different engine architectures. To use these models in a gradient-based engine design and optimisation tool, pyCycle, Kriging surrogate models are built using the Surrogate Modelling Toolbox, following the procedure described in Chapter 5. With this, the first research question is answered.

The second research question of this project concerns the effect of including emission constraints in optimisation; for this, a baseline engine model and an optimisation problem are defined in Chapter 6. Lastly, Chapter 7 contains the results of the optimisation cases described in Chapter 6, and an analysis of the effects of adding emissions constraints during preliminary engine design.

Future recommendations and limitations of this work are outlined in Chapter 9. To complement this thesis, Appendix A provides a sensitivity analysis of a Cantera reactor, to help understand the behaviour of the chemical reactor networks used. Appendix B provides an example case of a hydrogen turbofan with water recirculation, which was not part of the main set of results, but which shows the method being used in an unconventional engine and with an unconventional fuel at the same time.

The overall structure of the thesis is shown in Figure 1.1.

Project Background

Climate change is one of the biggest challenges for humankind in current times. In order to curve the current global warming trends, the scientific community and governmental agencies are working to reduce the impact that humans have on global warming. Aviation is one of the sectors that contributes to global warming through its emissions, and, as such, the aerospace community is involved in many different research efforts to reduce its global warming footprint. One of these research areas is propulsion, where research is focusing on making engines more efficient, and on finding less polluting energy sources.

The aim of the thesis is to propose a methodology for emissions prediction during preliminary gas turbine design which does not require detailed combustor geometry and allows for gradient-based optimisation of the engine. This methodology will be tested by optimising a conventional high bypass turbofan engine, the same architecture but with hydrogen, and a turbofan engine with water recirculation.

2.1. Climate Change and the Role of Aviation

At the current rate, the average Earth surface temperature is predicted to go up by more than 4 °C by 2100 [7] compared to pre-industrial temperatures (the average temperature in the second half of the 19th century), as can be seen on Figure 2.1. To try to avoid this, countries are trying to cut down on greenhouse gas (GHG) emissions. For example, the European Union aims to become carbon-neutral by 2050 [13]. As part of this target, the EU published the Flightpath 2050 goals for aviation, which are a 75% reduction in CO₂, a 90% reduction in NO_x emissions, and a 65% reduction in aircraft noise, all compared to novel aircraft from the year 2000 [14]. Additionally, 195 countries from the UN signed the Paris Agreement in 2015, of which the main goal is *"Holding the increase in the global average temperature to well below 2 °C above pre-industrial levels and pursuing efforts to limit the temperature increase to 1.5 °C above pre-industrial levels"* [15].

Aviation currently accounts for 11.6% of the CO₂ emissions from transportation [16], and 2.5% of the total CO₂ emissions [17]. To reduce this number, there is a need to find decarbonisation strategies. While increasing engine and aircraft efficiency is a way to reduce emissions, it is not enough to achieve the climate goals, and more radical innovations are needed. One possibility is finding alternative propulsion and power sources for aircraft that are less polluting.

Battery-powered aircraft are one option if electricity is produced using renewable sources, but with current and projected technologies, they will not be feasible in time to reach the aforementioned goals. With the available battery technology, electric aircraft may become a reality for general aviation and short-haul commercial aviation, but they will not be able to substitute current long-haul airliners [18].

Another option is the use of alternative fuels, such as hydrogen. Hydrogen is a carbon-free fuel that can be burnt in turbofans in a similar way in which kerosene is currently burnt. It can also be used in fuel cells to obtain electricity for electric motors. These two options are shown in Figure 2.2, obtained from Adler and Martins [8]. This research project will focus solely on hydrogen combustion in turbofans, so fuel cells will not be further considered here, but Adler and Martins [8] described fuel cells in detail.




Figure 2.1: Comparison of global warming forecasts depending on the preventive measures taken. Figure obtained from Ritchie, Rosado, and Roser [7].




Figure 2.2: Two most common configurations for aircraft with hydrogen propulsion. Image obtained from Adler and Martins [8].

Hydrogen turbofans can have a very similar architecture to current kerosene turbofans. However, since hydrogen properties differ significantly from kerosene, it cannot simply be used in current aircraft as a drop-in fuel. There are two main reasons for this; first and foremost, hydrogen needs high-pressure cryogenic tanks in the fuselage (it is not a liquid fuel that can be stored in wing tanks such as kerosene) [8]. Some examples of possible hydrogen tank configurations in conventional tube and wing aircraft are shown in Figure 2.3, obtained from Adler and Martins [8]. Secondly, the different thermodynamic and chemical properties of hydrogen mean that engine design needs some modifications, particularly in the combustor. However, with the necessary design changes hydrogen can be burnt in turbofans with similar architectures to current ones, which can be achieved with the current knowledge and technology [19], and thus hydrogen-powered long-haul aircraft can be a reality earlier than battery-powered ones.



Figure 2.3: Possible hydrogen tank configurations for tube+wing aircraft, where the darker shade of blue represents the passenger cabin. Image obtained from Adler and Martins [8].

Hydrogen has been considered as a potential aviation fuel for quite some time, and there have been many research efforts throughout different disciplines to study the use of hydrogen in aircraft. The next section provides a short summary of the potential of hydrogen in aviation and its properties.

2.2. Hydrogen in Aviation

Hydrogen is a carbon-free fuel with a high gravimetric energy density, meaning that it is lighter compared to other fuels (around 3 times lighter than kerosene for the same energy content). On the other hand, its volumetric energy density is low, so it would take more volume to store the same energy content than with other fuels (over 3 times more than kerosene for cryogenic liquid hydrogen). This can be seen in Figure 2.4 [9].




Figure 2.4: Comparison of volumetric and mass energy densities of different energy carriers for aviation. The two fuels marked in red are the two analysed in this project. The data in the figure was obtained from Yin and Gangoli Rao [9].

Unlike non-renewable fossil fuels, hydrogen has the potential to be a renewable energy source, in the sense that it can be obtained from water through electrolysis using green electricity (among other methods). However, please note that not all hydrogen production methods are renewable (e.g., hydrogen produced from electrolysis using non-renewable electricity, or hydrogen produced from methane).

2.2.1. Hydrogen Emissions versus Jet A Emissions

Current aeroengines run mostly on Jet A or Jet A1 (or comparable fuels), kerosene-type hydrocarbon fuels. Burning hydrocarbon fuels results in high CO₂ emissions. Due to incomplete combustion of the fuel, mostly in fuel-rich areas in the combustor, carbon monoxide (CO) and unburned hydrocarbons (UHC) are also present in the exhaust gas. Furthermore, soot is also present in the exhaust of kerosene-based engines.

Hydrogen fuel does not contain any carbon, meaning that all of the above emissions would be automatically eliminated through the use of hydrogen turbofans. However, burning hydrogen still produces emissions. NO_x emissions, which are caused by the oxidation of the nitrogen present in the atmosphere, are produced in any type of combustion in which nitrogen is present (either in the oxidiser or bound in the fuel composition), due to the high temperatures that are reached inside the combustor. The formation mechanisms for NO_x are further detailed in Section 3.2.

Moreover, hydrogen combustion results in high levels of water vapour, which may lead to different contrail formation properties compared to kerosene. Whether hydrogen turbofans would produce more or less contrails than kerosene engines is not yet clear, due to the complex formation mechanisms of contrails.

Hydrogen does not have any carbon emissions during flight, and it can be free from associated carbon emissions during production, as is the case with hydrogen produced through electrolysis using renewable electricity or nuclear energy (also known as green hydrogen and pink hydrogen, respectively). Nevertheless, several methods of producing hydrogen do have associated carbon emissions, such as electrolysis using electricity produced by fossil fuels, or hydrogen produced through steam-methane reforming. Thus, life cycle emissions should not be forgotten when considering an overall comparison between kerosene and hydrogen. Koroneos et al. performed a life cycle emissions comparison between kerosene and hydrogen, and while they

showed that hydrogen emissions are greatly dependent on the production method, kerosene emissions are higher in general when compared to hydrogen or other aviation fuels [20]. Ortuño et al. [21] performed a climate impact assessment comparing hydrogen and kerosene aircraft with similar performance capabilities, and they found that despite consuming more energy than kerosene aircraft, the atmospheric temperature response was 67% lower in hydrogen aircraft.

2.2.2. Logistics

Another one of the challenges of hydrogen is the infrastructure and logistics involved. Unlike kerosene, which is liquid at room temperature, hydrogen is a gas with a low volumetric energy density. For it to be used as an aircraft fuel, it needs to be transformed to cryogenic liquid hydrogen, or at least compressed to a high pressure. This means that all of the current infrastructure for kerosene would not be usable and needs to be replaced by completely new infrastructure [19], including storage tanks, pipelines, or pumps. Moreover, for hydrogen aircraft to be usable on a large scale, they need to have the possibility to refuel all around the world, and thus this new infrastructure needs to be deployed in a large number of airports; otherwise, the number of possible routes for hydrogen-powered aircraft would be limited.

The effect of hydrogen aircraft on ground operations is a current area of research, since the required changes of the operational procedures would be noticeable. If the required airport infrastructure for hydrogen aircraft is not ready to keep up with an increasing number of hydrogen aircraft, turnaround times could increase and lead to delays [22]. The refuelling process in particular for hydrogen aircraft would inherently be different than for kerosene aircraft, due to the nature of the fuels (pressurised gas or cryogenic liquid vs. ambient temperature liquid). Different sources report different expected refuelling times; Mangold et al. [23] report generally faster refuelling times for LH2 than for kerosene, whereas ten Damme [24] found that refuelling a hydrogen aircraft comparable to the Bombardier Q400 took 3 times longer than the original kerosene version of the Q400. Nevertheless, this longer refuel time does not necessarily imply a longer turnaround time if other tasks are parallelised [24].

2.3. Novel Turbofan Engine Technologies

To achieve the target reductions in aviation emissions mentioned in Section 2.1 and to decrease the fuel consumption of new aircraft, recent turbofan engines show architectural changes with respect to older generations. The new design trends, aimed at increasing the thermodynamic efficiency of the engines, include increases in the bypass ratio (BPR), the overall pressure ratio (OPR), and the turbine inlet temperature (TIT), as seen in Figure 2.5 [25]. However, these design trends encounter practical limits that hinder the increases in efficiency. For example, the TIT is limited by the turbines, since an increased TIT reduces their lifetime, and there are maximum temperatures that their materials can endure. Thus, increasing the TIT is associated with an increase in the required cooling flows to maintain the blade temperatures within the allowable range, which in turn implies a reduction in efficiency that can offset the benefits of increased TIT. Moreover, higher peak combustion temperatures result in significantly higher thermal NO_x emissions, as is explained in Section 3.2. Increasing the OPR must be done in conjunction with increasing the TIT in order to optimise the engine cycle [26], so its increase is associated with the same issues as increasing the TIT itself. Additionally, increasing the OPR means that the bleed flows for cooling are at a higher temperature, so a higher mass flow would be needed for turbine cooling, which is negative for the engine efficiency [26].

Due to these limits on the incremental innovation of turbofan engines, several new technologies are being investigated with the objective of bringing significant changes to gas turbine architecture that can bring the required fuel consumption and emission reductions. In this section, several of those technologies (i.e., the ones that will be considered in this research) are explained.

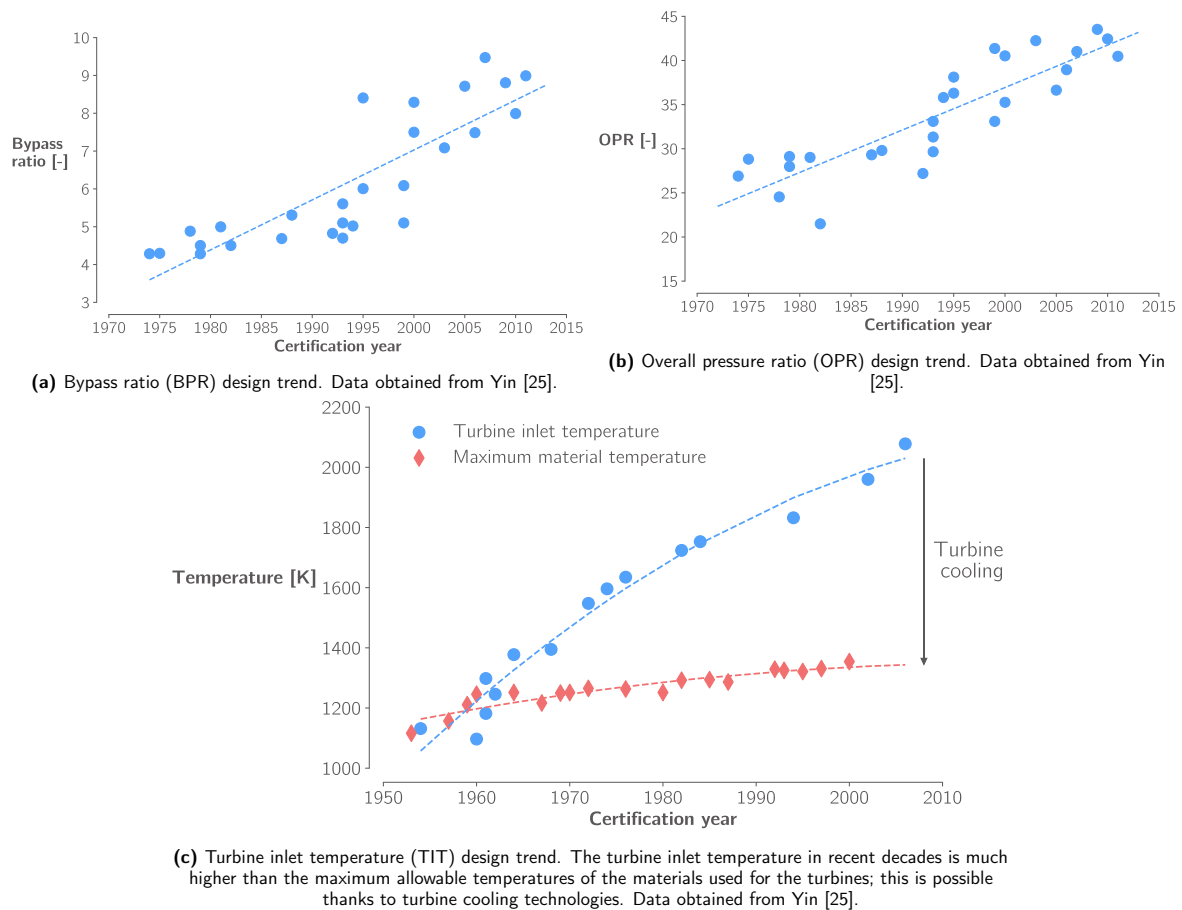


Figure 2.5: Engine design trends for (a) bypass ratio, (b) overall pressure ratio, and (c) turbine inlet temperature.

2.3.1. Challenges of Novel Architectures

Before diving into the aforementioned new technologies, it is important to understand the design challenges that novel architectures bring to (preliminary) design.

During preliminary design, there is no detailed geometry available for high fidelity simulations of the engine, so the use of simple algebraic correlations that only require high-level performance data from the engine is common. These correlations are usually based on the known performance of current engines, and are tuned based on reference performance data, and therefore they are useful for the preliminary design and analysis of engines that have similar technologies and architectures to the ones currently in use. However, the applicability of these correlations is limited to the range of the input parameters of the reference data, and to engines with similar technology. Therefore, they cannot be used to analyse engines that feature novel technologies, or that have performance metrics outside the bounds of the reference data (e.g., if a simple emissions correlation that takes the combustor inlet and outlet pressures and temperatures as inputs is only calibrated for a certain range of pressures and temperatures, it will not be applicable to a new engine that features significantly higher OPR and TIT, even if the basic architecture is the same).

Thus, one of the challenges of preliminary design is to be able to capture performance trends of novel architectures without relying on inapplicable correlations nor on unavailable detailed geometry. In this research, this challenge is the most obvious for the prediction of emissions of the proposed new architecture. This specific topic is covered in more detail in Chapter 4.

2.3.2. Water Injection

Atma et al. [27] analysed the performance effects of recirculating water on a turbofan based on the N+3 engine. This was done by extracting exhaust water and injecting it before the HPC. The study compared four configurations: Jet A without water recirculation (baseline), Jet A with water recirculation, H₂ without water recirculation, and H₂ with water recirculation. They concluded that a hydrogen engine with water recirculation has a higher efficiency than the baseline Jet A engine, and that the high amount of exhaust water vapour on hydrogen engines makes the concept viable (water injection in Jet A engines would require onboard water tanks that may completely offset the benefits with the extra weight of the water).

Water injection has several potential benefits, such as decrease in thrust specific fuel consumption (TSFC), reduction in NO_x emissions, or higher lifetime of the hot components of the engine [10, 28]. However, these benefits and the overall effects on performance differ depending on the injection location within the engine (directly in the combustor, before the HPC, before the LPC, etc.), and whether the water is injected as steam or in liquid form (in which case the size of the droplets also has an effect on performance). This section covers the different types of water injection and its effects on performance, as well as the side effects it has on required turbine cooling. A sample diagram of a turbofan with water recirculation can be seen in Figure 2.6.

On top of the performance benefits, Mourouzidis et al. conducted a study on water injection including the possible economic repercussion, taking British Airways as the reference airline in the case study [29]. They found that the technology could bring savings close to £600,000 annually, and recoup the investment in around 7.5 years.




Figure 2.6: Sample water injection diagram with water misting before the HPC and LPC. Image obtained from Daggett et al. [10].

Injection Location

The location of the water injection influences the effect that the added water has on cycle performance. The main distinct locations are 1) before the low pressure compressor (LPC), 2) before the high pressure compressor (HPC), and 3) into the combustor. The effects that each of these injection locations has on the engine performance were analysed by NASA, Boeing, and Rolls-Royce [10]. They found that injecting atomised water before the LPC results in higher available thrust and efficiency, resulting in either an increase in thrust for the same fuel consumption, or a lower fuel consumption for the same thrust level. Both of these options translate into a lower TSFC. Furthermore, a reduction in NO_x is observed due to the lower combustor inlet and exit temperatures. For water injection before the HPC, they found that while the same performance and emission improvement trends were observed, the gains were lowered for the same amount of water compared to injection before the LPC. The potential benefit of injecting water before the HPC rather than before the LPC is the lower risk of droplets freezing in subzero conditions. Lastly, for water injection directly into the combustor, they report higher reductions in NO_x for the same amount of water, at the cost of a decrease in engine efficiency and performance. The reduction in flame temperature due to the water injection causes thermal losses in the combustor for all three cases, but these losses are offset by performance improvements due to increases compressor mass flow in the cases in which water is injected

before the compressors.

Daggett et al. [10] explain in their report that the droplet size is relevant to the performance of water injection before the LPC, with bigger water droplets (above 10 μm) evaporating less promptly and thus being centrifuged outwards in the LPC, reaching and cooling the case and affecting blade tip clearance. Moreover, they report that the distance between the water atomisation and the start of the LPC has an effect on droplet evaporation. These effects combined result in a need of having water droplets of at most 5-10 μm in order to achieve the performance and emission benefits stated above.

Another effect that Daggett et al. [10] report is that, for the same thrust levels, shaft speeds and bypass mass flow are both reduced due to the increased core mass flow, resulting in a small reduction in engine noise.

Experimental data from water injection before the compressor [30] and in the combustor [31] shows agreeing trends with the results presented above.

Steam versus Liquid Water

Walsh and Fletcher [32] describe similar trends for liquid water injection on Chapter 12 of their Gas Turbine Performance book. Additionally, they provide information on the comparison between liquid water and steam. They explain that steam has similar trends to water injection, except that TSFC increases with steam injection into the combustor rather than decreases. However, it must be noted that in their explanation they assume that steam is generated with waste heat from the exhaust, which account for part of the performance gains. In a turbofan with water recirculation, exhaust heat would likely be the source of thermal energy to convert the collected liquid water to steam (should steam injection be chosen over liquid water injection). However, an additional heat exchanger in the exhaust would introduce pressure losses before the core nozzle, and add extra weight, so a trade off would need to be performed between water and steam injection.

2.4. Baseline Project

This project aims to build on a previous project conducted at the MDO Lab of the University of Michigan on gradient-based optimisation of a turbofan with water recirculation based on the NASA N+3 engine. The results of that project were published in a paper titled *Comparing Hydrogen and Jet A for an N+3 Turbofan with Water Recirculation using Gradient-Based Optimization* [27].

The project compared two different configurations of the turbofan: with and without water recirculation, each tested with Jet A and H_2 as fuels, making a total of four different cases.

The water injection was modelled with two ad-hoc components, a water injector and a water extractor. The water extractor removes a certain percentage of the water content of the exhaust gases, and the injector adds that corresponding mass flow to the air upstream of the high pressure compressor (HPC) as water vapour. Figure 2.7 shows the turbofan configuration and the position of these components. One limitation of this set-up is that the change of phase in water is not considered; the water content is extracted and injected as vapour.

The optimisation problem was a gradient-based multi-design point (MDP) optimisation with four flight conditions: top of climb (on-design), rolling take-off, static sea level, and cruise. The optimisation algorithm used was SNOPT (Sparse Nonlinear OPTimizer) [6]. The objective was to minimise fuel flow during cruise while maintaining 5800 lbf (25.8 kN) of thrust at top of climb, and having a fan diameter smaller than 100 inches (2.54 m).

The project showed that water recirculation led to an improvement of approximately 5% in thrust-specific energy consumption for both Jet A and hydrogen (this metric is further explained in Section 3.1). The hydrogen engine with water recirculation is approximately 4% more efficient than the 'dry' Jet A engine. Hydrogen combustion produces higher quantities of water in the exhaust, meaning that there is more water available for recirculation and therefore the concept is more feasible in hydrogen engines, and designing




Figure 2.7: Configuration of the two-spool turbofan with water recirculation.

an efficient-enough water condenser can therefore be less challenging than for the Jet A counterpart. In contrast, for effective water injection in Jet A engines, carrying extra water in tanks might be necessary, with the weight penalty that it carries. Thus, the project concluded that water recirculation is a promising option for hydrogen engines that should be further explored.

2.5. Framework and Tools

The framework under which this research is being carried is presented in this section. This includes an introduction to multidisciplinary design and optimisation (MDO), as well as an introduction to the tools that will constitute the basis of this research.

2.5.1. What is MDO?

Aircraft design is subject to the continuous goal of improving efficiency and reducing fuel consumption. To keep improving aircraft designs with respect to previous generations, making aircraft that are just good enough is not sufficient, and trying to iterate the design by hand will not lead to the necessary efficiency gains any more. Moreover, aircraft design is a group effort of several disciplines that are inherently coupled to each other, and what would be optimal for one of them might not even meet the requirements for another (e.g., a wing optimised solely considering aerodynamics might not be thick enough to house a structure that can safely handle all required loads, or might not have sufficient volume to store the required fuel). Thus, the design of aircraft is a complex balance between all disciplines, which can be addressed through multidisciplinary design and optimisation (MDO) [33].

MDO consists in the simultaneous numerical optimisation of several coupled disciplines of a certain system. The benefits of MDO is that it can lead to a final design with better performance than through the conventional iterative design process, while also reducing the costs of the design cycle and the uncertainty in the design [33].

It is important to clarify that while the explanation above is focused around overall aircraft design, MDO is applicable to many other disciplines outside of aerospace, and can also be applied to a certain individual system of the aircraft (i.e., the engine, which will be the case in this research project).

2.5.2. OpenMDAO

OpenMDAO [34] is a Python MDO framework that is designed to perform gradient-based optimisation using analytical derivatives that are explicitly provided by the user when the components of the optimisation problem are defined. It is based on the MAUD (modular analysis and unified derivatives) architecture from Hwang and Martins [35]. The tool also offers the possibility of automatically computing the gradients of components for which the analytical derivatives are not available, through either finite difference (FD) or complex step (CS) differentiation [36].

2.5.3. Cycle Analysis: pyCycle

The cycle analysis tool that will be used in this thesis is called pyCycle [4], a Python-based tool built on top of OpenMDAO with the objective of performing gradient based optimisation of aeroengines. This tool is based on NPSS (Numerical Propulsion System Simulation) by NASA [37], but features hand-derived analytical derivatives for all of the available components, which allows for much faster optimisation runtimes. In the paper presenting the tool, Hendricks and Gray verified pyCycle against NPSS and showed an agreement within 0.03% of most thermodynamic properties [4].

Hendricks and Gray [4] also performed a comparison between the optimisation times of pyCycle and NPSS, using a model of the N+3 turbofan by NASA [5] as the optimisation baseline. They observed that the optimisation times with pyCycle were up to 22 times faster than NPSS, when computing the NPSS derivatives with a monolithic approach using finite difference of the entire code (treating it as one black box), and 7 to 15 times faster when compared to semi-analytically-computed NPSS derivatives (the partial derivatives were computed using finite difference, but the total derivatives were computed using the direct or adjoint methods) [4].

In pyCycle, several components are offered to allow the construction of arbitrary engine architectures. Moreover, since it is an open source Python package, it allows the user to build their own components if necessary. For example, Atma, Lamkin, and Martins [27]. created their own components for water extraction and injection to analyse the water recirculation effects in the cycle, since those components were not available originally in pyCycle.

Emissions Modelling

To reduce the emissions of an engine, the first step is understand them and to be able to predict them, to later be able to adjust the engine design to minimise those emissions. This chapter begins with background theory on NO_x emissions and their formation mechanism, followed by emission modelling techniques. This knowledge is then applied to the creation of the two emissions models with the tool Cantera [38], one for Jet A and one for hydrogen. These models are integrated with pyCycle in Chapter 5, and are used in this project to analyse and optimise the performance of a turbofan engine with emission constraints.

3.1. Hydrogen Combustion

Before diving into the details of NO_x formation and emission modelling techniques, it is important to first understand that hydrogen combustion differs from Jet A combustion. These differences mean that combustors built for Jet A are different than those built for hydrogen, which in turn means that different models will be needed to analyse each of the two fuels.

The stoichiometric fuel-to-air ratio (FAR), defined as the ratio of fuel and air at which all of the fuel and all of the oxidiser are consumed, is 1/14.4 for kerosene, compared to 1/34 for hydrogen. However, in practical applications, fuel is not burnt at stoichiometric conditions (the actual used FAR depends on a combination of factors, such as emissions or flame stability). The ratio between the actual fuel-to-air ratio and the stoichiometric fuel-to-air ratio is called the equivalence ratio. Fuels cannot be burnt at any equivalence ratio, they must be burnt within certain limits in order to have stable, controlled combustion. Burning with too low equivalence ratio would lead to lean blowout, and burning with too high equivalence ratio would lead to rich blowout. These maximum and minimum limits are different for hydrogen and kerosene; for hydrogen, the flammability range is at equivalence ratios between 0.2 and 0.65, whereas for kerosene the range is 0.6 to 1.15, approximately [39]. These limits can be seen in Figure 3.1, obtained from Adler and Martins [8]. The fact that hydrogen can burn at lower temperatures means that the TIT can be reduced, which can result in the doubling of the turbine life [8, 40].

Another important difference between hydrogen and kerosene are the flame speeds. Hydrogen has a flame speed approximately 4 times higher than kerosene [39]. This is a very important consideration when designing the combustor, since the risk of flashback is significantly higher for premixed hydrogen combustors due to the higher flame speed. Together with the higher flame speed, the higher diffusivity of hydrogen leads to higher reaction rates and thus to lower residence times compared to Jet A, which in turn lead to lower thermal NO_x emissions [41].

On the other hand, hydrogen has a higher adiabatic flame temperature when compared to Jet A. The gas mixture in the combustor can have hot pockets in which the fuel burns at (close to) stoichiometric ratio due to uneven mixing, and as a result of its higher adiabatic flame temperature, the hot pockets for hydrogen can be hotter than for Jet A, leading to higher NO_x production in those areas. Thus, to achieve low NO_x combustion of hydrogen, good mixing of the fuel and oxidiser is paramount, but, as stated in the previous paragraph, premixing comes with the risk of flashback [8]. Achieving good mixing while preventing flashback




Figure 3.1: Comparison of the flammability limits of hydrogen and kerosene. Figure obtained from Adler and Martins [8].

is one of the challenges of combustor design for hydrogen gas turbines.

All these differences can be summarised by saying that hydrogen and Jet A combustion are different, and so are the combustors used for each type of fuel. Thus, separate models need to be used to predict the emissions of hydrogen and Jet A engines.

3.2. NOx Emissions

Emissions of nitrogen oxides (NOx) are third on the list of aviation pollutants. Not only are they a GHG, but they are also toxic. Unlike the water vapour that leads to contrails, or the carbon that is needed for CO₂ formation, NOx emissions are a result of the oxidation of the nitrogen present in the atmosphere (although some of the NOx formation pathways use fuel-bound nitrogen or require hydrocarbon radicals, as will be explained below). Thus, changing the fuel from Jet A to H₂ will not eliminate NOx emissions, and thus accurately predicting NOx emissions in hydrogen-burning turbofans becomes one of the key aspects needed for their implementation.

NOx comprises nitric oxide (NO) and nitrogen dioxide (NO₂). Both of these compounds are produced during the combustion process in turbofans, and big part of these emissions are NO molecules that later oxidise to NO₂ (further described in Section 3.2.7). There are several formation mechanisms that result in these NOx molecules. Thermal NOx is caused by the oxidation of nitrogen under high temperatures, and prompt NOx is caused by hydrocarbon radicals interacting with nitrogen molecules. Other molecules such as N₂O or NNH also contribute to NOx formation. An overview of these pathways can be seen in Figure 3.2 [11], and they are explained in more detail in the following subsections.





Figure 3.2: Diagram representing the different pathways for NOx formation: thermal NOx (green), prompt NOx (red), N₂O-intermediate NOx (blue), NNH-pathway NOx (yellow), and fuel NOx (purple). Image obtained from Iavarone et al. [11].

3.2.1. Thermal NO_x

The thermal NO_x pathway, also called the Zeldovich mechanism after Yakov Zeldovich, is a chemical route which describes the oxidation of nitrogen into nitric oxide under high temperatures [42]. The production rate increases exponentially with temperature, so limiting the peak temperatures in the combustor, and residence times at such temperatures, is key to reduce thermal NO_x formation, which is the leading formation mechanism in conventional combustors. The reactions from this mechanism are shown in Reactions 3.1-3.3.



3.2.2. Prompt NO_x

Prompt NO_x, first described by Fenimore [43], is a mechanism of NO_x formation that results from the interaction of hydrocarbon radicals and nitrogen. The reactions proposed by Fenimore were the following:



However, this original mechanism faced criticism due to the fact that it does not conserve electron spin [44, 45], although it still was widely accepted and seemed to be in agreement with chemical kinetic studies performed in the following years [45]. Some decades later, Moskaleva and Lin [46] studied the kinetics of the reaction and proposed an alternative reaction to Reaction 3.4 that would conserve electron spin, with cyanonitrene radical instead of hydrogen cyanide as an intermediate species, shown in Reaction 3.6. This new reaction was able to converge the different theories and previous chemical kinetic experiments of prompt NO_x formation [45].



3.2.3. N₂O-intermediate NO_x

A new route for NO_x formation with nitrous oxide (N₂O) as an intermediate species was proposed by Malte and Pratt [47]. Production of NO_x through this route is less significant than other pathways at most conditions, but it becomes more notable at lower temperatures in lean conditions [48].







In Reaction 3.7, M is a third-body (which can be any molecule) that absorbs the energy of the formation of N₂O and carries it away as kinetic energy.

3.2.4. NNH Pathway

The importance of the role of NNH on nitrogen reactions was first raised by Miller et al. [49] after an investigation of the reactions involved on "Thermal DeNO_x" (reduction of NO through ammonia [50–52]). It was later Bozzelli and Dean [53] who proposed a set of reactions describing the role of diazenyl NNH on NO_x formation. They proposed the following reactions for NNH formation and subsequent NO formation:



In their paper, they also proposed two other reactions that produce species that might lead to further NO_x production through other pathways, although Reaction 3.13 has a rate constant an order of magnitude higher than these other two NNH + O reactions [53].



3.2.5. Fuel NO_x

Hydrocarbon fuels can have bound nitrogen atoms in their composition. During the combustion process, these atoms break loose and transform to NO_x, and may account for a significant part of NO_x emissions. The difference between fuel NO_x and prompt NO_x is that prompt NO_x is formed from nitrogen in the air in the presence of hydrocarbon radicals produced as a result of hydrocarbon fuel combustion, whereas fuel NO_x is formed from nitrogen that is chemically bound to the hydrocarbon fuel before its combustion.

The main pathways for fuel nitrogen conversion to nitric oxide are shown in Reactions 3.16 and 3.17 [54–56]; please note that these are high-level paths and not the detailed reactions.



3.2.6. NO_x Reburning

In an architecture with two combustors the NO_x molecules produced in the first combustor will enter the second combustor before reaching the atmosphere. These NO_x molecules may get partially eliminated

through a process known as NO_x reburning, in which nitric oxide (NO) interacts with hydrocarbon species to form hydrogen cyanide (HCN), or with NH_{*i*}, as shown in Reactions 3.18 and 3.19 [57].



The aforementioned reaction with hydrocarbons is prevalent on rich conditions; however, under lean conditions the hydrocarbon species react with oxygen to form carbon monoxide (CO), as shown in Reaction 3.20 [58].



Nicolle and Dagaut [59] analysed the role of NO_x reburning in flameless combustion and they reported that the NO-HCN reaction is specially relevant before autoignition of the mixture. Reburning does not strictly need two separate combustors; for example, in a study of NO_x emissions of an ITB with flameless combustion, Perpignan et al. described reburning due to recirculation in the ITB when there was no presence of NO_x in the oxidiser (the ITB was experimentally tested on its own without an upstream main combustor) [60].

3.2.7. NO Oxidation to NO₂

As it can be seen from Reactions 3.1-3.17, the formation of NO_x in the combustion process leads mostly to NO, whereas the definition of NO_x comprises both NO and NO₂. One reason for this is that, at high temperatures, NO₂ tends to dissociate into NO according to Equation 3.21 [48].



This does not mean that there are no NO₂ emissions resulting from gas turbines. Further down the engine, with lower temperatures, NO may convert to NO₂ under certain conditions. The presence of unburnt hydrocarbons can significantly increase this conversion [48], and as such, it can be expected than in hydrogen engines this conversion is less prevalent. However, NO can still be converted to NO₂ without the presence of hydrocarbons, through the hydroperoxyl radical HO₂, as shown in Equation 3.22 [48].



Once released into the atmosphere, nitric oxide (NO) can also oxidise to nitrogen dioxide (NO₂) [12].

3.2.8. Effects of Steam on the Chemical Kinetics of NO_x Formation

The addition of water or steam to the cycle has effects on the chemical kinetics in the combustor. Zhao et al. [61] report that for a certain flame temperature an increase in steam decreases the CH concentration and as a result the prompt NO_x formation. Similarly, they report that the decrease in CH leads to a decrease in the N radical from HCN. Moreover, they report that the increase in steam results in a higher concentration of OH, because the steam is decomposed into OH, and additionally it slows down the reactions that consume OH. This increase in OH leads to more NO_x from the OH reactions, but the decrease in the N radical leads to an overall reduction in NO_x formation due to steam [61]. The NNH pathway is also reduced under an increased presence of steam [59].

3.3. Semi-Empirical Correlations

One of the most extended methods to estimate NO_x emissions during preliminary design is the use of simple algebraic correlations that are tuned based on current and past engines. These correlations take simple inputs that are available during preliminary design, such as the inlet and outlet pressures and temperatures in the combustor, and return an estimated value of the emissions index for NO_x (EINO_x).

One of the most widely used correlations is the P₃-T₃ method, shown in Equation 3.23 [62, as cited in, 63], which, as the name suggests, takes the inlet pressure and temperature of the combustor (commonly referred to as p₃ and T₃) as input variables. EINO_x is the emission index of nitrogen oxides (g of NO_x/kg of fuel), FAR is the fuel to air ratio, and h is the humidity. The subscripts *SL* and *FL* stand for sea level and flight level, respectively. This method can predict the emissions at a certain flight level when the sea level emissions are known. The coefficients *n* and *m* can be tuned to specific engines if the data is available, or *n*=0.4 and *m*=0 can be used as a best-estimate for generic engines.

$$\text{EINO}_{x_{FL}} = \text{EINO}_{x_{SL}} \left(\frac{p_{3,FL}}{p_{3,SL}} \right)^n \left(\frac{\text{FAR}_{FL}}{\text{FAR}_{SL}} \right)^m e^H \quad (3.23)$$

with $H = 19(h_{SL} - h_{FL})$

This method is effective for conventional turbofan configurations that burn Jet A. However, this method is not tuned for engines with configurations that are significantly different to those of current engines (i.e., it would not be wise to use it for an engine with an ITB or with water recirculation), nor for engines with different fuels (e.g., hydrogen). Moreover, the method requires the emissions at sea level, which would not be available nor easy to estimate for an unconventional configuration.

Viola et al. [64] proposed an extension of the P₃-T₃ method to make it work for hydrogen, and tuned it with data from air turborockets. As shown in Equation 3.24, the Mach number (*M*) and the Damköhler number (*Da*, defined as residence time over ignition delay time) are included in the equation to include the effects of compressibility in different conditions, and to account for the residence time of the fuel in the combustor, respectively.

$$\text{EINO}_{x_{FL}} = a \text{EINO}_{x_{SL}} \left(\frac{p_{3,FL}}{p_{3,SL}} \right)^b \left(\frac{\text{FAR}_{FL}}{\text{FAR}_{SL}} \right)^c M^d \left(\frac{\text{Da}_{FL}}{\text{Da}_{SL}} \right)^f e^H \quad (3.24)$$

This extended method follows the same principle of the original P₃-T₃ method, and therefore it still requires the EINO_x at sea level. Moreover, while this solves one of the shortcomings of the original method (inapplicability when using hydrogen instead of Jet A), the method is still invalid for engines with water recirculation.

Since there are currently no flying turbofan engines that are fuelled on hydrogen (and very few engines with water injection), there are currently no semi-empirical correlations that are applicable to all the configurations that will be studied in this thesis. Therefore, semi-empirical correlations are too limiting for the goal of this project, and another approach should be used to estimate engine emissions.

3.4. Chemical Reactor Networks

A chemical reactor network (CRN) is a modelling technique in which a system is represented by one or more components (reactors) in order to analyse the chemical kinetics of the flow inside that system. These components are called (chemical) reactors, and CRNs are composed of an arbitrary number of them. An aeroengine combustor can be represented by a CRN in order to analyse the combustion process inside it, and thus predict emissions without relying in empirical correlations that are purely based on past engines.

There are different types of reactors, and the ones that will be considered for this application are the perfectly

stirred reactor (PSR) and the plug-flow reactor (PFR). These two types of reactors are further explained in the following subsection.

3.4.1. Cantera

Cantera [38] is an open-source Python package for chemical kinetics and thermodynamics with which chemical reactor networks can be built. It is the tool that is used to implement the CRNs described in Chapter 4.

To construct the desired network architecture, Cantera offers different reactors (described in Section 3.4.2) and reservoirs (components with infinite volume and constant properties that work as sources or sinks for the flows), as well as auxiliary components, such as mass flow controllers or pressure controllers, to connect the reactors and reservoirs.

The tool offers the possibility of selecting any desired reaction mechanism from an input file, meaning that all of the reaction mechanisms presented in Section 3.5 are compatible with Cantera.

3.4.2. Types of Reactors

A CRN is composed of several building blocks known as reactors, in which the chemical reactions take place. Depending on the characteristics of the flow and the assumptions used, there are different types of reactors that can be used. A CRN can be formed by a single reactor, multiple reactors of the same type, or even multiple reactors of different types, in order to best approximate the flow conditions of the system that is to be analysed. The two types of reactors used in the CRNs described in Chapter 4 are the perfectly-stirred reactor (PSR), sometimes also called well-stirred reactor or continuously-stirred reactor; and the plug-flow reactor (PFR). This two types of reactors are explained below.

Perfectly Stirred Reactor

Perfectly stirred reactors (PSRs) are a type of ideal reactors in which perfect mixing of the flow is achieved within the control volume, and in which the outlet properties of the gas are equal to those in the control volume of the reactor. Figure 3.3 shows a diagram of a generic PSR, obtained from Turns [12].




Figure 3.3: Representation of a perfectly stirred reactor (PSR). Image obtained from Turns [12].

Equations 3.25-3.28 show the conservation equations that define a PSR [12]. Equation 3.25 defines the conservation of mass for each species i , where the \dot{m} terms represent the mass flow going in and out of the reactor, and $\dot{m}_i'''V$ represents the production or elimination rate of each species (due to chemical reactions). Equation 3.26 is a rewritten version of Equation 3.25, using the definitions $\dot{m} = \dot{\omega}_i MW_i$ (relation between mass generation rate and net production rate) and $\dot{m}_i = \dot{m}_{tot} Y_i$ (the mass flow of species i is equal to the total mass flow times the mass fraction of species i).

$$\frac{dm_{i,cv}}{dt} = \dot{m}_{i,in} - \dot{m}_{i,out} + \dot{m}_i''' V \quad (3.25)$$

$$\dot{m}(Y_{i,in} - Y_{i,out}) + \dot{\omega}_i MW_i V = 0 \quad (3.26)$$

In these equations, m is the mass [kg], \dot{m} is the mass flow [kg/s], \dot{m}_i''' is the mass production rate per volume [kg/s/m³], and V is the volume [m³]. Y_i is the mass fraction of species i [kg of species i /kg of mixture], $\dot{\omega}$ is the species production rate [kmol/s/m³], and MW is the molar mass [kg/kmol].

Equation 3.27 is the conservation of energy equation, where $h_i(T)$ is the enthalpy of species i , defined in Equation 3.28.

$$\dot{Q} = \dot{m} \left(\sum_{i=1}^N Y_{i,out} h_i(T) - \sum_{i=1}^N Y_{i,in} h_i(T_{in}) \right) \quad (3.27)$$

$$h_i(T) = h_{f,i}^0 + \int_{T_{ref}}^T c_{p,i} dT \quad (3.28)$$

In these equations, \dot{Q} is the heat transfer rate [W], h is the enthalpy [J/kg], h^0 is the enthalpy of formation [J/kg], c_p is the specific heat at constant pressure [J/kg/K], and T is the temperature [K].

Plug-Flow Reactor

Plug-flow reactors (PFRs) are reactors in which the properties of the flow may change along the axial direction, but are constant throughout the entire cross-section for any certain axial position. Any kind of axial diffusion is neglected [12]. A schematic is shown in Figure 3.4.




Figure 3.4: Representation of a plug-flow reactor (PFR). Image obtained from Turns [12].

The system of ordinary differential equations (ODEs) that define the plug-flow reactor is shown in Equations 3.29–3.31 [12].

$$\frac{d\rho}{dx} = \frac{\left(1 - \frac{R_u}{c_p MW_{mix}}\right) \rho^2 v_x^2 \left(\frac{1}{A} \frac{dA}{dx}\right) + \frac{\rho R_u}{v_x c_p MW_{mix}} \sum_{i=1}^N MW_i \dot{\omega}_i \left(h_i - \frac{MW_{mix}}{MW_i} c_p T\right)}{P \left(1 + \frac{v_x^2}{c_p T}\right) - \rho v_x^2} \quad (3.29)$$

$$\frac{dT}{dx} = \frac{v_x^2}{\rho c_p} \frac{d\rho}{dx} + \frac{v_x^2}{c_p} \left(\frac{1}{A} \frac{dA}{dx}\right) - \frac{1}{v_x \rho c_p} \sum_{i=1}^N h_i \dot{\omega}_i MW_i \quad (3.30)$$

$$\frac{dY_i}{dx} = \frac{\dot{\omega}_i MW_i}{\rho v_x} \quad (3.31)$$

The equations represent the change in density (ρ in $[\text{kg}/\text{m}^3]$), temperature (T in $[\text{K}]$), and mass fraction (Y). In these equations R_u is the universal gas constant ($\text{J}/\text{kmol}/\text{K}$), v_x is the speed in the axial direction of the reactor (m/s), A is the cross-sectional area of the reactor (m^2), and P is the pressure (Pa). The other parameters follow the same definition as for the PSR.

This type of reactor is not directly available in Cantera. Thus, to model it for the CRN described in Section 4.3, the plug-flow region is discretised as a series of PSRs, where each PSR effectively represents one control volume cv from Figure 3.4.

3.5. Reaction Mechanisms

CRN analysis relies on reaction mechanisms to compute the actual chemistry in the reactors, usually as an external input file to the programme. These mechanisms are a simplified set of reactions and rate constants that describe the process that wants to be analysed. Since it would be prohibitive to analyse every single reaction that exists, these mechanisms are a simplification of the real chemistry that includes only the reactions and species of interest. For example, a mechanism made for hydrocarbon combustion will include many species of hydrocarbons, hydrocarbon radicals, CO_2 , CO , etc. However, a mechanism made to analyse hydrogen combustion will most likely not include any species containing carbon. Furthermore, some hydrogen mechanisms include all relevant reactions for NO_x formation, whereas other hydrogen mechanisms are further simplified and do not include any species with nitrogen and cannot be used for NO_x emissions prediction.

For this work, two mechanisms were selected, one for kerosene engines and another one for H_2 engines.

3.5.1. Luche's Mechanism for Jet A

In order to model Jet A combustion, a mechanism that contains the required hydrocarbon species must be used. Moreover, not only the hydrocarbon combustion process needs to be included in the mechanism, but also the relevant reactions (and species) needed for the prediction of emissions must be present in the mechanism. Luche presented a reduced mechanism for Jet A combustion that includes NO_x formation and which consists of 91 species and 991 reactions, together with two other simpler versions featuring 91 species and 694 reactions, and 89 species and 680 reactions, respectively [65, 66]. This reduced mechanism is based on a more detailed mechanism originally developed by Dagaut et al. [67, 68], in which kerosene is represented as 74% n-decane ($\text{nC}_{10}\text{H}_{22}$), 15% n-propylbenzene (phC_3H_7), and 11% n-propylcyclohexane ($\text{cyC}_9\text{H}_{18}$), and in which combustion is modelled through 225 species and 1800 reversible reactions [66].

This mechanism has been used in the literature to model kerosene combustion in aeroengine combustors when NO_x emissions were relevant [1, 2].

3.5.2. Meng et al.'s Mechanism for H_2 with NO_x pathways

Mechanisms like Luche's mechanism, or the popular GRI-Mech 3.0 [69] (designed for natural gas combustion), contain atomic and molecular hydrogen in their species and reactions, which means they allow the modelling of hydrogen combustion. However, in pure $\text{H}_2 + \text{air}$ combustion, where there are no hydrocarbons present, the prompt and fuel pathways described in Section 3.2 do not contribute to the formation of NO_x , since they depend on those hydrocarbons. Thus, to model hydrogen combustion, using a different mechanism with its reaction rates tailored to pure hydrogen combustion can improve the accuracy of NO_x estimation.

Meng et al. [70] studied the role of HNNO in NO_x formation during combustion, and suggest that a new NO_x pathway, mediated by HNNO , can have a prominent role in NO_x formation under certain conditions. Consequently, they developed a mechanism for NO_x emissions based on their findings, which is based on a mechanism by Glarborg et al. [71], with a HNNO submodel appended to it.

This model is the one that was used in this project for the hydrogen combustion models. Other H_2 combustion

models, such as the one proposed by Capurso et al. [72] based on a premixed, swirl-stabilised hydrogen flame, have been tested, but the mechanism by Meng et al. showed the best match against the reference data.

Chemical Reactor Networks for Jet A and H₂ Combustors

Chapter 3 explained the fundamental information required to understand NO_x emissions in combustion, and the available modelling techniques that can be used to predict them. Chemical reactor networks are the best fit to the requirements of this project, and thus are the chosen approach to do emissions prediction in this project. This chapter describes the two chemical reactor networks used in this project, one for Jet A combustion and one for H₂ combustion. The network architectures are first described, followed by the verification of the implementation, and the description of any modifications made to adapt the networks and their inputs to the requirements of this project.

4.1. Calculation of EINO_x

The models presented in this chapter are focused on predicting NO_x emissions. The metric used for this is the emissions index of NO_x (EINO_x). This is calculated using Equation 4.1, where Y_{NO_x} is the mass fraction of NO + NO₂, and FAR is the fuel-to-air ratio. The units are grams of NO_x per kilogram of fuel (g/kg).

$$\text{EINO}_x = 10^3 \frac{1 + \text{FAR}}{\text{FAR}} Y_{\text{NO}_x} \quad (4.1)$$

4.2. Jet A Combustor Model

A common approach to modelling combustors with chemical reactor networks is creating an ad-hoc model for the combustor that needs to be analysed, for which detailed geometry information and flow field data from CFD simulations is used. However, this approach is not useful for early design stages in which no geometry is available. Villette et al. [1] proposed a simplified CRN architecture to analyse conventional Jet A aeroengine combustors without needing detailed geometry. The architecture of the network is shown in Figure 4.1.

The CRN consist of three perfectly stirred reactors, each corresponding to the primary, secondary, and dilution zones of the combustor, which are the three regions in which the combustion in conventional combustors is usually divided. These reactors are fed by 2 reservoirs containing the air and fuel, respectively. The fuel and airflow into the reactors is controller by mass flow controllers in Cantera, and the flow out of each reactor is controlled by a valve. All the fuel is injected into the primary zone, whereas the air is distributed across the three zones.

To achieve ignition of the mixture an artificial igniter is used. A short pulse of hydrogen atoms is injected into the primary zone, where these atoms trigger the combustion reaction in the primary zone. The ignited mixture propagates downstream through the reactors, achieving combustion in the entire network. Once

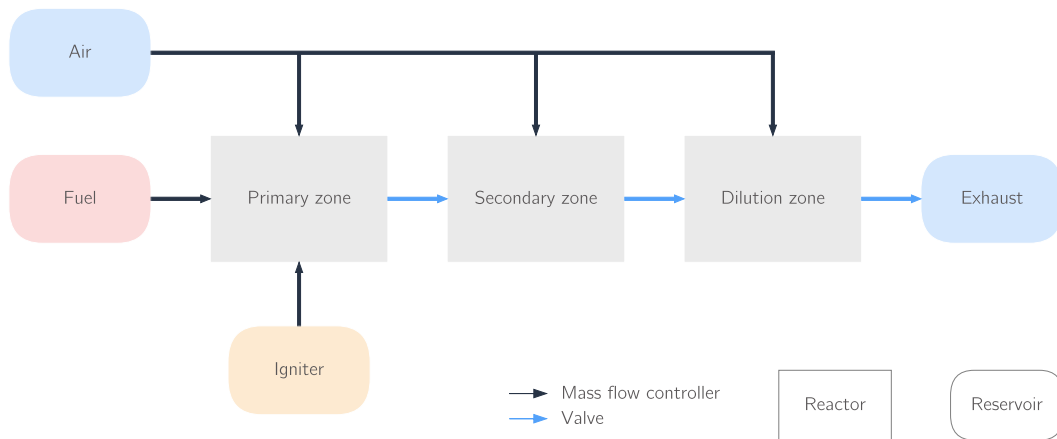


Figure 4.1: The chosen network for conventional Jet A combustors, based on the architecture proposed by Villette et al. [1], which consists of three reactors connected by valves, fed by reservoirs containing the oxidiser (air) and the fuel (Jet A), and features an igniter reservoir from which a pulse of hydrogen atoms flows to start the combustion reaction in the primary zone.

combustion is initiated, the combustion reaction is self sustained; the igniter pulse only acts transiently at the beginning of the simulation, and its effects disappear once converged to steady state. This means that as long as combustion is successfully initiated, the ignition pulse does not affect the steady state results.

There is one difference in implementation between their proposed architecture and the CRN used in this work; the pressure loss in their architecture was implemented through a single valve behind the dilution zone. In this work, the pressure loss has been distributed evenly across the three zones (with three separate valves) to represent the pressure loss along the combustor, while maintaining the same total pressure loss as in their model.

The model takes the inputs shown in Table 4.1. The total mass flow into the combustor is specified, and the percentage of that mass flow that goes into the primary and secondary zone is also defined; the remaining air goes into the dilution zone. Similarly, the total length of the combustor, and the individual lengths of the primary and secondary zones, are defined; the dilution zone length is calculated as the remaining length.

Table 4.1: Inputs parameters for the Jet A model.

Parameter	Units	Description
T_3	K	Temperature at the combustor inlet
p_3	Pa	Pressure at the combustor inlet
\dot{m}_{air}	kg/s	Mass flow of air into the combustor
FAR	—	Fuel-to-air ratio
dPqP	%	Pressure loss along the combustor
AR_{PZ}	%	Percentage of the total air injected into the primary zone
AR_{SZ}	%	Percentage of the total air injected into the secondary zone
A	m^2	Cross-sectional area of the combustor
L	m	Total length of the combustor
L_{PZ}	%	Percentage of the total length corresponding to the primary zone
L_{SZ}	%	Percentage of the total length corresponding to the secondary zone
Igniter parameters	—	Parameters of the igniter pulse

The reaction mechanism used in this model is Luche's mechanism described in Section 3.5.1. The air is

modelled as 21% oxygen and 79% nitrogen by mole fraction, and the fuel is modelled as 74% $nC_{10}H_{22}$, 11% cyC_9H_{18} , and 15% toluene.

4.2.1. Verification

Villette et al. only provide temperature results for one of the analysed cases, the CFM56 engine in cruise condition. They provide the temperature profile for each of the three reactors. Using the input parameters provided in their paper [1], which are listed in Table 4.2, the same case was simulated to verify the CRN implementation. The results are shown in Figure 4.2, where the results show a very good match (maximum relative error of 0.64%) for the three zones. Small deviations are expected despite using the same CRN architecture, since the input values reported in their paper, are probably rounded (this is assumed based on the fact that reported percentages that should add up to 100% only add up to 99.99%, for example). Moreover, the temperatures in their paper were provided in a graph, so they had to be digitised, which may have led to small errors.

Table 4.2: Inputs parameters for the CFM56-7B27 in cruise condition [1].

Parameter	Value	Units
T_3	687	K
p_3	9.67e5	Pa
\dot{m}_{air}	17.35	kg/s
FAR	0.022	—
dPqP	5.02	%
AR_{PZ}	30.82	%
AR_{SZ}	23.02	%
A	0.16	m^2
L	0.178	m
L_{PZ}	29.9	%
L_{SZ}	33.6	%




Figure 4.2: The reproduced model matches the temperature profile of Villette et al.'s architecture [1] within a maximum relative error of 0.6% across the three combustor zones.

4.2.2. Model Tuning

Villette et al. tested their model in three different engines: The CFM56-7B27, the CFM LEAP-1A26, and the Rolls-Royce Trent 772. For each of the engines they provided different air ratios and length ratios of the combustor zones. While this leads to more freedom in tuning each of the engine models, it represents a difficulty in terms of using the model in a generic manner in pyCycle, since pyCycle does not provide these ratios. Thus, in order to have a model that represents all combustors, the air and length ratios were tuned to obtain the best possible common values for the three analysed engines.

The tuning objective was the minimisation of the error between the EINO_x predicted by the model, and the EINO_x reported in the ICAO Database [73]. For each of the three engines there were four tuning points: take-off, climb, approach, and idle, leading to a total of 12 tuning points. The tuning was performed using the absolute error, the relative error, and the root mean square error (RMSE). After comparison of the results, it was determined that the RMSE tuning resulted in the best matching of the trends, so it was selected as the tuning objective. The tuning function is shown in Equation 4.2.

$$J = \sqrt{\frac{\sum_n^{12} (\text{EINO}_{x_n, \text{CRN}} - \text{EINO}_{x_n, \text{ICAO}})^2}{12}} \quad (4.2)$$

The tuning was performed using SciPy's minimize function with the Nelder-Mead method [74]. The results of the tuning are shown in Table 4.3.

Table 4.3: Tuned results of the air ratios in the primary and secondary zones, and the length ratios of the primary and secondary zones.

Parameter	Tuned value	Unit
AR _{PZ}	28.03	%
AR _{SZ}	26.72	%
L _{PZ}	6.01	%
L _{SZ}	45.31	%

Using these tuned parameters, the NO_x emission predictions of the three engines across the four operating points were recalculated. Figure 4.3 shows the comparison between the tuned results (blue), the target ICAO data [73] (black), and the data reported by Villette et al. [1] (red). The latter was digitised from the figures provided in their paper. The objective of the tuning and the model is to capture the NO_x emission trends, more than accurately matching each point, and the results show matching trends with respect to the reference data. While tuning each individual engine or operating condition would provide a more accurate match at each point (as is the case with the data from Villette et al.), that would defeat the objective of having a generic model that can predict trends across different engines and operating conditions.

4.2.3. Extension to Water Recirculation

The original model is built to analyse conventional engines in which water injection is not present. To extend the capability of the model to analyse water injection, an extra input parameter was added, the water fraction in the inlet air. The 'dry' air composition of the baseline model is 21% oxygen and 79% nitrogen. The 'wet' model adds steam to the primary zone reactor from a separate reservoir containing 100% water, with the water mass flow becoming an additional input function to the reactor network.

4.3. H₂ Combustor Model

As described in Section 3.1, hydrogen combustion differs from kerosene combustion. Thus, combustors made for hydrogen combustion are also different from those built for Jet A fuel. For this reason, a separate

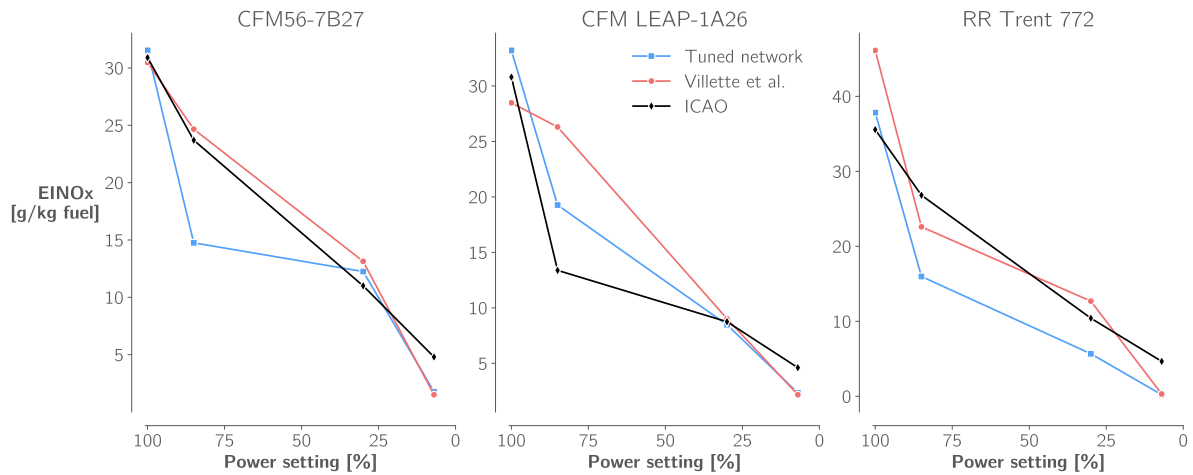


Figure 4.3: The tuned data has air ratios and zone lengths common among the three engines, whereas Villette et al.'s data has different values for each engine. With the common parameters the NO_x emission trends are captured despite the error in the actual values being greater than the original for some of the cases. The data from Villette et al. was digitised from their figures [1].

combustor architecture is used to simulate hydrogen NO_x production. The chosen architecture, shown in Figure 4.4, is based on the work of Talboom [2].

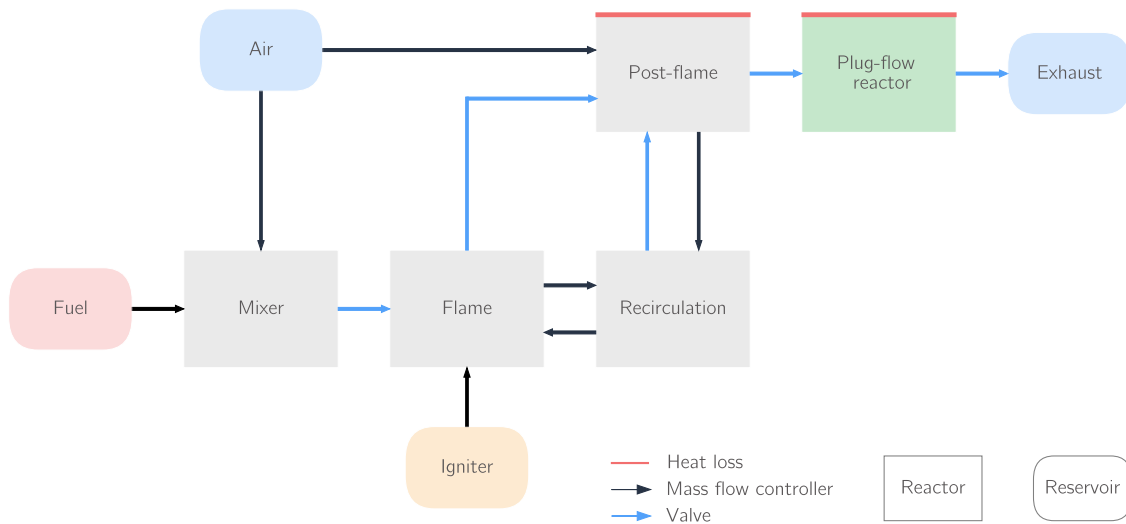


Figure 4.4: The architecture proposed by Talboom [2] features 4 perfectly stirred reactors followed by a plug-flow reactor. There is recirculation between the flame, post-flame, and recirculation reactors. The reactors are fed by reservoirs containing the air and fuel, and ignition is modelled by a pulse of hydrogen atoms injected into the flame reactor.

The architecture consists of four perfectly stirred reactors, followed by a plug-flow reactor. The reactors are fed by air and fuel reservoirs, and ignition is achieved by a pulse of atomic hydrogen, in the same way that was described in Section 4.2.

This model is more complex than the one used for Jet A combustion. The first main difference is that the air does not simply flow downstream, but there is recirculation among the flame, recirculation, and post-flame reactors. The second difference is that the model features a plug-flow reactor, placed downstream of the post-flame reactor. This plug-flow reactor cannot be simulated directly in Cantera with a specific function, so it is modelled as a series of PSRs following the approach described in Section B.2. In contrast with the Jet A model, ignition here does not happen in the first (mixer) reactor, which is only there to simulate the mixing of air and fuel before the flame region, it happens in the second (flame) reactor.

All the fuel is injected into the mixer reactor, and the air is injected into the mixer and post-flame reactors. The air flowing into the post-flame region does not represent additional air being injected downstream in the real combustor, it is done to capture the effects of unmixedness between the air and the fuel (since PSRs assume perfect mixing, adding all of the air and fuel into a PSR would assume perfect mixing in the combustor). The amount of air going into the post-flame region is controlled by the Bypass Oxidiser Factor (BOF), and it is set to be 5% across all conditions that are presented in Section 4.3.1 [2].

The recirculation between the flame, recirculation, and post-flame reactors is dictated by a mix of mass flow controllers and valves. The mass flow controllers dictate the flow between the reactors, and the valves allow flow to scape to maintain constant pressure and equilibrium along the reactors. There are three parameters that dictate the flow: the Recirculation Mass Flow Ratio (RMFR), the Recirculation Inlet Distribution (RID), and the Recirculation Outlet Distribution (ROD). The first one, the RMFR, defines the amount of recirculation mass flow as a fraction of the total inlet mass flow, and is estimated to be $1.2 - \varphi$ [2].

The recirculation mass flow is the combined mass flow that goes into the recirculation reactor from the flame and post-flame reactors. The RID defines what percentage of the recirculation mass flow comes from the flame reactor, and how much comes from the post-flame reactor. Talboom provides a value of 35% of the flow coming from the flame reactor and 65% coming from the post-flame reactor [2]. This is dictated by mass flow controllers in the reactor network.

The mass flow entering the recirculation reactor also needs to leave to maintain equilibrium, and it does so by going back to the flame and post-flame reactors. The parameter that controls how much of the recirculation mass flow goes back to the flame reactor is the ROD. Talboom refers a ROD value of 20% [2], meaning that 20% of the flow goes back to the flame reactor (through a mass flow controller). The remaining flow (80%) goes back into the post-flame region, but to maintain equilibrium in the network this is not enforced by a mass flow controller. Instead, a valve connects the two reactors, allowing for equilibrium of the simulation during transient pressure and mass flow fluctuations, while still converging to the desired 80% output at steady state.

This combustor architecture was used to represent an experimental combustor that was 0.297 metres long, and had a circular cross section with a diameter of 0.104 metres. The combustion regions inside the combustor (and thus the corresponding reactors) are defined as described in Figure 4.5.

The dimensions described in the figure were defined as follows; these relations, given in millimetres, were obtained from Talboom's thesis [2], and are calculated from experimental flow field data.

- r_1 is calculated as: $37 - 26\varphi$.
- R is fixed to 52 mm.
- x_1 is fixed to 28.81 mm. Talboom provides a relation that makes this parameter dependent on the fuel momentum ratio for momentum ratios higher than 2. However, since the injection velocities are not know, this value has been left as constant. This approximation does not seem to negatively affect the results, as seen in the comparisons from Section 4.3.1.
- x_2 is fixed to 34 mm.
- x_3 is calculated as: $17 + 64\varphi$.
- x_4 is fixed to 200 mm.
- L is fixed to 297 mm.

The inputs to the model are described in Table 4.4. These inputs are the inputs of the baseline network used in the verification cases. In Section 4.3.2, some modifications are done to the network in order to adapt it to a full size combustor under operation conditions, which expands the list of inputs.

4.3.1. Verification

To verify the implementation of the results, three separate tests were considered. The first test compares the outlet temperatures of the combustor in adiabatic conditions, the second test compares the same cases

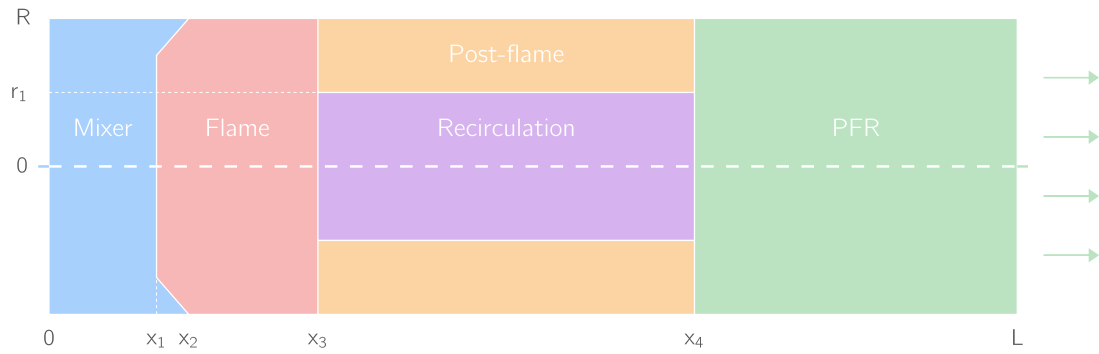


Figure 4.5: Streamwise cross-section of the combustor. The vertical dimension corresponds to the diameter of the cylindrical combustor, and the horizontal direction is the length. The flow goes from left to right. The figure and the dimension definitions are based on the description from Talboom [2].

Table 4.4: Inputs parameters for the baseline H₂ model.

Parameter	Units	Description
T_3	K	Temperature at the combustor inlet
φ	—	Equivalence ratio
ζ	—	Fraction of heat lost to the environment
Igniter parameters	—	Parameters of the igniter pulse

as the first test, but including heat loss to the environment; and the last test is a comparison of the NO_x emissions across a different range of inlet temperatures and equivalence ratios.

For the adiabatic case, the network was simulated with an inlet temperature of 453 K for equivalence ratios ranging from 0.4 to 0.8. The results of this calculation are shown in Figure 4.6, which shows the good match between the network built in this project and the original results by Talboom, with the maximum relative error being below 0.3%.

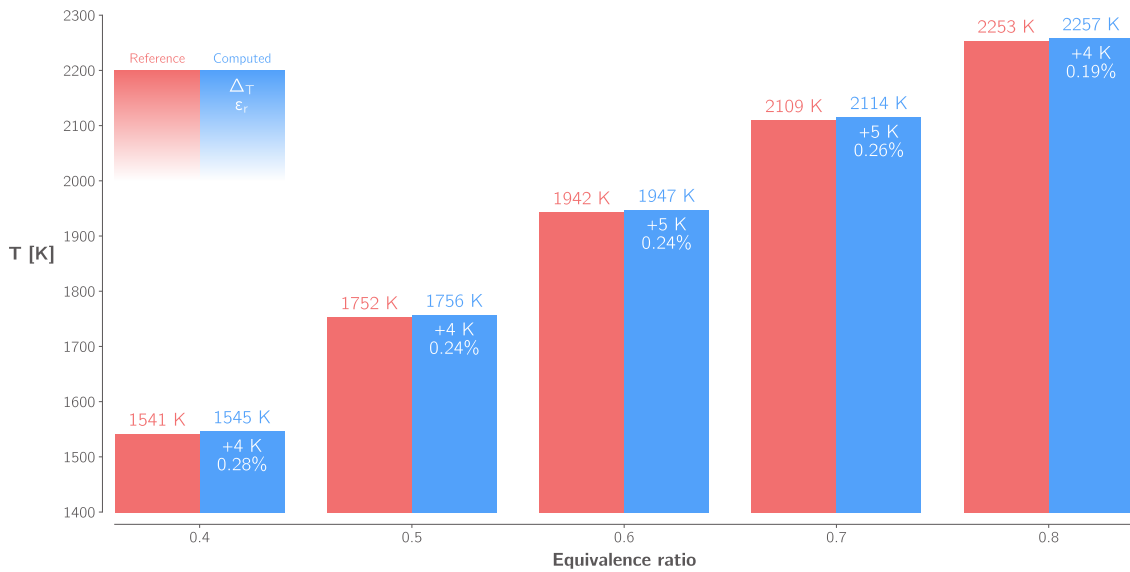


Figure 4.6: The results of the adiabatic verification cases show very good matching across different equivalence ratios, with the maximum error being within 0.3%. The simulated temperature is consistently higher across all points by 4/5 K.

The same conditions (inlet temperature of 453 K and equivalence ratios from 0.4 to 0.8) were simulated including heat loss to the environment. The heat loss was modelled as a fraction of the total heat produced in the combustor, and the percentages were set to be proportional to the equivalence ratio according to Equation 4.3, the relation used in the reference results [2].

$$\zeta = 0.23\phi + 0.23 \quad (4.3)$$

The heat loss in Cantera is set as a heat flux, which is calculated from the fuel flow and its lower heating value (LHV) (representing the total heat generated), and then multiplied with the heat loss fraction described above, as shown in Equation 4.4. Half of this heat loss happens in the post-flame region, and the other half in the PFR.

$$q = \zeta \dot{m}_f \text{LHV}_{H_2} \quad (4.4)$$

The results (in blue) were compared against the reference CRN results reported by Talboom (in red) and the experimental results (in black), with the comparison shown in Figure 4.7. The absolute and relative errors shown in the blue bars represent the errors of the computed results compared to the experimental results. The match is good across all equivalence ratios, with the maximum relative error being 2.52% in the case with the highest equivalence ratio (0.9).

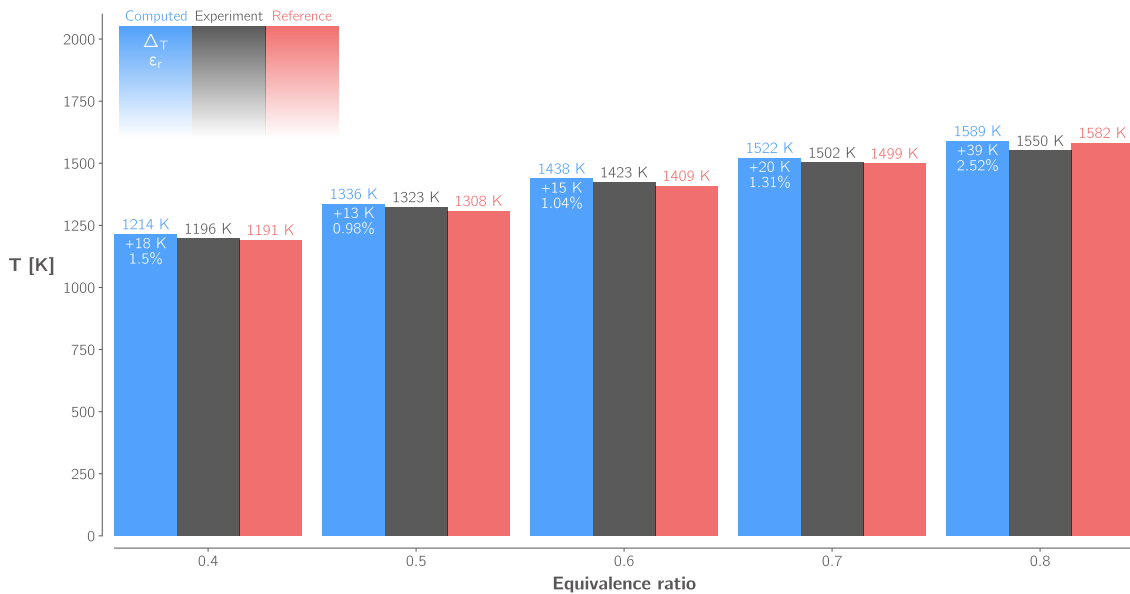


Figure 4.7: The results of the verification cases with heat loss show good a good overall match, with a maximum relative error of 2.52% at the highest equivalence ratio. The absolute and relative errors are of the computed values (blue) with respect to the experimental data (black). The red values represent the results obtained by Talboom's CRN [2].

For the last test, the NO_x emissions were compared for four different inlet temperatures: 313, 453, 623, and 703 K. For each temperature, different equivalence ratios were analysed. The computed results were compared against the reference results computed by Talboom, and the experimental results from the AHEAD project as reported in Talboom's work [2]. The results show good match in terms of both trends and actual values for all temperatures and equivalence ratios with the exception of one single case; the predicted NO_x emissions at an equivalence ratio of 0.9 (the highest analysed) with an inlet temperature of 313 K (the lowest) are underpredicted by close to 35%. However, in this project the model will be used in operating conditions with equivalence ratios closer to 0.3 and higher operating temperatures, as will be seen in Chapter 7, and therefore this one-off discrepancy is considered acceptable.

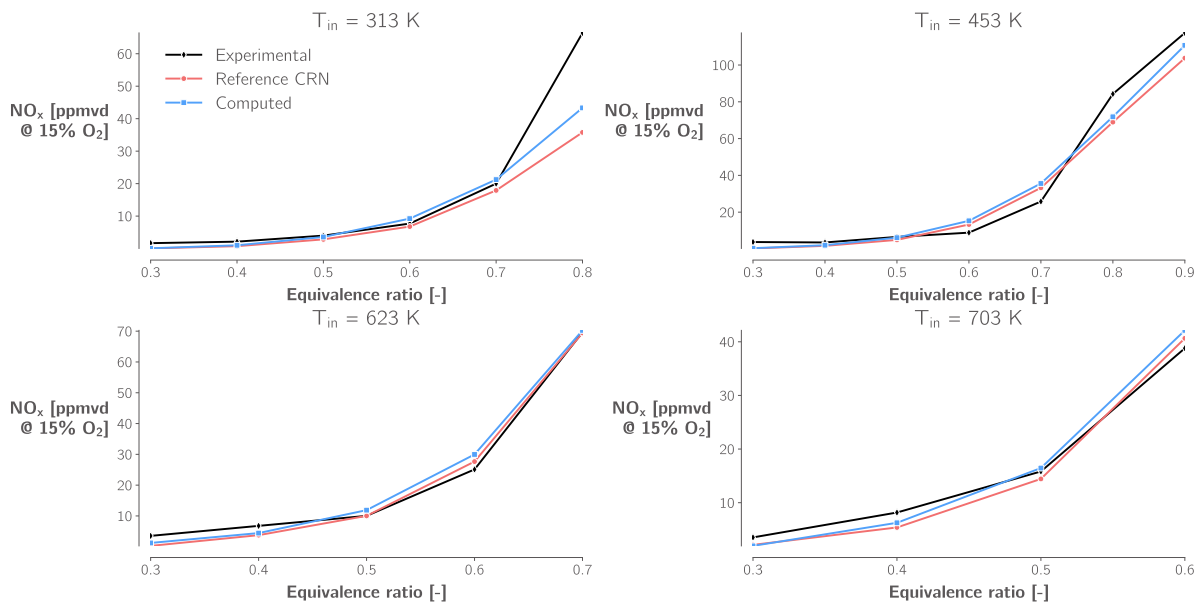


Figure 4.8: The comparison of the obtained NO_x trends show very good agreement with both the experimental (black) and reference (red) results from Talboom [2]. A noticeable underprediction of the NO_x levels occurs at the case with the lowest inlet temperature (313 K) combined with the highest equivalence ratio (0.9).

4.3.2. Extension to Full Size Combustor

The architecture presented above was tuned to reproduce the results of an experimental campaign conducted in a small combustor operated at ambient pressure. The results of the experiment were matched accurately, but the model cannot be used as is. Several modifications need to be performed to extend the model to an actual engine combustor. Talboom provides a series of modifications to achieve this extensions [2], which were also applied here with some minor modifications to match the engine architectures used in this project.

The first modification corresponds to the value used for the bypass oxidizer factor, which was set to 5% in the experimental cases. For the full combustor, Talboom suggests a value of 15%, which is tuned for an equivalence ratio of 0.3. As will be seen in Chapter 7, this equivalence ratio of 0.3 is very close to the operating conditions of the engines analysed in this project.

The second modification is the removal of the heat loss to the environment. The AHEAD experiment was carried out in an isolated combustor that lost a significant amount of heat to the environment, so this heat loss had to be modelled to reproduced those results. However, the full combustor is considered adiabatic, so no heat loss is applied in its corresponding model.

The third modification is also related to the heat management of the combustor. While all the inflow air in the experiment was injected directly in the combustor inlet, in a real combustor, which operates at higher temperatures, the combustor walls need to be cooled with air, which also enters the combustor and dilutes the flow. Following the recommendations, the amount of air entering the combustor (to the mixer reactor) was set to 80%, with the remaining 20% being used as dilution air. This dilution air is fed to the PFR region. As mentioned previously, the PFR is not modelled directly in Cantera, so it is reproduced as a series of small PSRs. All PSRs have the same volume, which is the total volume of the PFR region divided by the number of PSRs used to model it. In this work, 500 PSRs have been used to model the PFR region. The dilution air is divided evenly among all these PSRs, with the total amount of air injected adding up to the 20% of dilution air. The schematics of this set-up are shown in Figure 4.9.

To simulate the PFR region, a new network, separate from the PFR region, is built. This is done for simplicity in the construction of the network, but does not affect the results compared to simulation the PSR and PFR regions in the same network, since there is no upstream effect from the PFR region into the PSR one. In this network, the first PSR is fed gas with the composition and thermodynamic properties of the converged




Figure 4.9: A plug-flow reactor is modelled as a series of perfectly stirred reactors, all having the same size. The dilution air is fed evenly to all PSRs. Instead of being connected to the rest of the network, the PFR region is modelled separately, where a reservoir with the properties of the post-flame region feeds the first PSR.

post-flame reactor from the PFR region, where the mass flow rate is set to the outlet mass flow rate of the post-flame region. This first PSR is converged to steady state, and its converged composition and thermodynamic properties are then fed to the second reactor. This process is then repeated until all reactors are converged.

Lastly, the dimensions of the combustor need to be scaled to its full size. Talboom suggests a scaling of 3.1 to the volumes of the reactor, while maintaining all the volume ratios of the individual reactors constant. However, for the optimisation, the combustor dimensions should be able to vary, and to adapt to the cross-sectional area inputs from the optimiser. Therefore, the dimensional scaling here was done as follows: first, the original combustor dimensions were computed using the equivalence ratio relations provided for the original combustor; secondly, using the cross-sectional area from pyCycle, the actual combustor diameter is computed (assuming a circular cross-section); lastly, all dimensions of the combustor are scaled by the ratio of $D_{full}/D_{original}$.

The reactor network architecture of the full combustor with these modifications applied is shown in Figure 4.10.

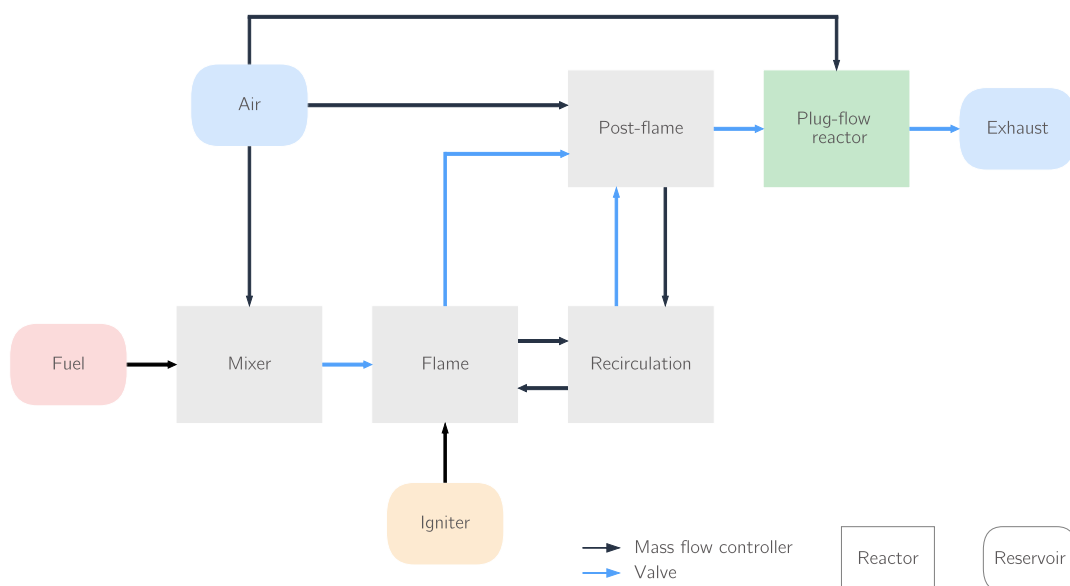


Figure 4.10: The full architecture features no heat loss to the environment, and a dilution stream connected to the PFR.

5

pyCycle Integration through a Surrogate Model

The objective of this work is to obtain a methodology to include emission predictions in gradient-based optimisation during preliminary design of novel architectures. The first part of the objective, estimating emissions of novel architectures during preliminary design, has been achieved in Chapter 4. However, the models cannot be used as-is in pyCycle, since Cantera does not offer gradient information of the reactor network calculations.

The simplest approach to integrate these models in pyCycle would be to wrap them in an OpenMDAO explicit component, and calculate the gradients with finite difference. However, this approach has several drawbacks.

The first one is that the Cantera models are not robust enough to handle the optimisation. During the optimisation, the optimiser might test points that deviate significantly from both the starting point and the optimum point, and, since the Cantera models may fail for different reasons (e.g., if the mass flow is too high with respect to the igniter pulse, successful ignition might not be achieved, or if the input parameters deviate significantly from the expected range the simulation might not converge).

The second problem is that finite differences are not efficient. Running a single point of the Cantera models is not computationally insignificant, and the model would need to be probed once for each input variable at each iteration. Given the high number of input variables and the cost of running the model, this approach would increase the computational time of the optimisation significantly.

Lastly, not using a methodology that can handle gradients intrinsically would be against the objective of this project.

To circumvent these problems, the approach taken was to create a surrogate model for each of the networks presented in Chapter 4. The process followed, and the tools used to create these surrogate models, will be discussed in this chapter, preceded by some background information on surrogate modelling.

5.1. Surrogate Modelling

Surrogate models are approximations of functions or models that are constructed from fitting input and output data of the original model. Their main benefit is that, once constructed, they can be faster to query than the original models that they substitute.

In the case at hand, a surrogate model would substitute the entire Cantera model, providing a model that is much faster to compute than the reactor network, and that avoids the issues of the simulation not converging. Moreover, there are tools available which provided surrogate modelling capabilities with gradient information. Thus, surrogate models solve the three main issues that the Cantera models presented for optimisation.

To construct the surrogate model, the Cantera models are computed at a number of different input conditions, and the Cantera output (in this case, the EINO_x) is recorded. The surrogate model takes the inputs and outputs from Cantera (and nothing else), and fits the data to provide an approximation of the Cantera model.

5.1.1. Sampling: Latin Hypercube Sampling

The first step to generate the model with the procedure above is to have the different input conditions from which the model will be sampled. There are several sampling techniques that are used in the literature to build surrogate models, such as random sampling or full factorial sampling. In this work, the sampling technique chosen is Latin hypercube sampling (LHS).

LHS is a more efficient version of random sampling that avoids the clustering issues that random sampling might present, and requires a lower number of points to have a satisfactory distribution. The main difference between LHS and random sampling is that in LHS the points are not selected independently of each other. The space is divided into n intervals for each variable, where n is the number of sampling points, and the points are placed such that there is only one point in each interval. Additional considerations can be present depending on the specific implementation of the LHS method [33], for example centring the points within the intervals, simply ensuring that there is only one point per interval but choosing the actual points randomly, or maximising the minimum distance between points. The latter approach, usually referred to as 'maximin', is amongst the most popular in LHS, since it allows to cover the design space with fewer points than if the points are selected randomly, and it is the chosen approach in this work.

5.1.2. Modelling: Kriging

Once the sampling points are selected, the output of the Cantera models is computed for each point. With the inputs and outputs, the surrogate model can be trained. However, for this a model needs to be chosen first. Here again there are several options in the literature to choose from, such as least-squares, radial basis functions, or polynomial approximations.

In this work, the chosen modelling method is Kriging. Unlike other modelling techniques that attempt to fit the sampling data by tuning a function, Kriging does this by tuning a statistical model that aims to capture the behaviour of the function. The basic formulation of Kriging is shown in Equation 5.1 [33].

$$F(x) = \mu(x) + Z(x) \quad (5.1)$$

The first part of the equation, $\mu(x)$, represents the trend of the function or its behaviour, and $Z(x)$ represents the output of a stochastic process. There are several implementations for both $\mu(x)$ and $Z(x)$.

A common approach is to let $\mu(x)$ have a constant value, and use $Z(x)$ to handle the model behaviour. This approach is called *ordinary Kriging* [33]. This was the approach used in this work.

The stochastic part $Z(x)$ represents the error from the trend function (in this case a constant value) and is used to represent the behaviour of the function. If two sampling points are very close together, the Z values for those points are expected to be close, and the further away the sampling points are, the less that their Z values are correlated. In a Kriging model this correlation is handled by a kernel (covariance function) that dictates the correlation between the Z values corresponding to two points, as described in Equation 5.2 [33].

$$K(x_i, x_j) = \text{corr}(Z(x_i) Z(x_j)) \quad (5.2)$$

There are several kernel functions available in the literature for Kriging models. The most common one and the one used in this work is the Gaussian kernel, also called squared exponential kernel, which is defined by Equation 5.3 [3, 33, 75], where the subscript l represents the component of the vector \mathbf{x} , and θ_l is the

hyperparameter in the l dimension that controls the behaviour of the kernel.

$$K(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = \exp\left(-\sum_{l=1} \theta_l |\mathbf{x}_l^{(i)} - \mathbf{x}_l^{(j)}|^2\right) \quad (5.3)$$

5.1.3. Surrogate Modelling Toolbox

The tool chosen to implement the surrogate models is the Surrogate Modelling Toolbox (SMT) [3]. The SMT is an open-source Python package that provides sampling and surrogate model methods, and is focused on handling derivatives. Both Latin Hypercube Sampling and the Kriging method described above are available in the SMT.

5.2. Variable Scaling

The variables in Table 5.1 differ by several orders of magnitude, from $\mathcal{O}(10^6)$ for the pressure to $\mathcal{O}(10^{-2})$ for the FAR or the cross-sectional area. Using these inputs directly in the Kriging model would result in the variables with higher order of magnitude dominating the calculations, and the performance of the model. To avoid this, the inputs were scaled using standard deviation scaling. This results in every input having a distribution with the same mean (0), and with the same standard deviation. This scaling is performed with Equation 5.4, where \mathbf{x}_l is the vector of sampling points for the input variable l , $\mu_l(\mathbf{x})$ is the mean of that vector, and σ_l is its standard deviation.

$$\mathbf{x}_{l_{scaled}} = \frac{\mathbf{x}_l - \mu_l(\mathbf{x})}{\sigma_l} \quad (5.4)$$

This allows the Kriging model calculations to consider the effect of each variable during the fitting process without giving more priority to inputs with bigger magnitudes or a bigger range.

5.3. Jet A Surrogate Models

The first step to generate the surrogate models is to select the points that will be used to sample the chemical reactor network and to train the model. To do this, the bounds for the inputs variables, and the number of sampling point per variable need to be chosen.

The bounds were selected such that all the possible combinations analysed by the optimiser during the engine optimisation fall within the sampling space. The limits are shown in Table 5.1.

Table 5.1: Bounds of the sampling intervals for each of the input functions of the Jet A surrogate model.

Parameter	Unit	Lower bound	Upper bound
T_3	K	600	1000
p_3	Pa	10e5	22e5
\dot{m}_{air}	kg/s	9	20
FAR	—	0.017	0.03
A	m ²	0.02	0.08
L	m	0.1	0.25
\dot{m}_{water}	kg/s	0	1.5

Two separate surrogate models were constructed, one for the model including steam as an input variable (i.e., including all variables in Table 5.1), and another one for the 'dry' engine (using rows 1-7 in Table 5.1

as the input variables).

5.3.1. Combustor Length Calculation

One of the inputs to the surrogate model is the combustor length. While pyCycle provides the cross-sectional area of the combustor (and other components), it does not provide combustor length information, nor a volume to infer it from. Thus, a method to size the combustor length was added to the model.

The sizing approach was based on the constraint reported by Walsh that the residence time in the combustor should be at least 3 ms to allow for complete combustion [32]. The residence time can be estimated from the combustor length and the flow speed inside the combustor, as described in Equation 5.5.

$$t_{res} = \frac{L_{comb}}{v_{comb}} \quad (5.5)$$

The speed inside the combustor is not directly reported by pyCycle, but the reference Mach number is. Using the combustor inlet temperature T_3 to calculate the speed of sound in the combustor, the speed can be obtained from Equation 5.6. The heat capacity ratio γ for the combustion gases is set to 1.33, the specific gas constant R is 287.05 J/kg/K, and the Mach number inside the combustor is provided by pyCycle.

$$v_{comb} = M\sqrt{\gamma RT_3} \quad (5.6)$$

This calculation was added as an OpenMDAO explicit component to the pyCycle engine model. The combustor was sized at the rolling take-off (RTO) point (the multipoint design set-up is described in Chapter 6), since it was the case with the highest T_3 and mass flow, and hence the limiting case for residence time.

5.3.2. Cross-validation and Number of Sampling Points

Afzal et al. [76] studied the optimal number of sampling points when using Latin hypercube sampling for different surrogate models, including Kriging. They found that using a number of sampling points equal to 15 times the number of variables was a good starting point when using LHS, independently of the type of surrogate model chosen, but they also reported that Kriging showed the biggest improvement with an increase in sampling points among all analysed models. At the same time, it is important to avoid overfitting the model and creating unrealistic oscillations. To evaluate the effect of the number of sampling points on the error of this model, a k-fold cross-validation study was performed.

In a k-fold cross validation study, the available sampling points are divided in k subsets, then the surrogate model is trained on $k-1$ subsets, and tested with the remaining subset. This is repeated k times such that every combination of training and testing sets is analysed. This allows to evaluate the error of the model in a more 'data-efficient' manner, since all sampling points are used for training and testing at some point, so no subset needs to be excluded and dedicated exclusively for testing. In this work, the k value used was 5, i.e., the dataset was divided in 5 subsets, where 4 are used for training and 1 for testing, a 80-20 split.

This was then repeated for a different number of sampling points, from 6 up to 35 sampling points times the number of design variables. This was done by first generating a sampling set of $35 \times n$ points (e.g., 210 total points for the dry Jet A model that has 6 input variables) using Latin hypercube sampling, then taking subsets of n sampling points, and lastly performing the 5-fold cross validation study for each of this subsets. The error metric used in the results from this section is the root mean square error (RMSE).

Model without water recirculation

The results of this study for the dry Jet A model shown in Figure 5.1, where the effect of the number of sampling points on the error can be seen.

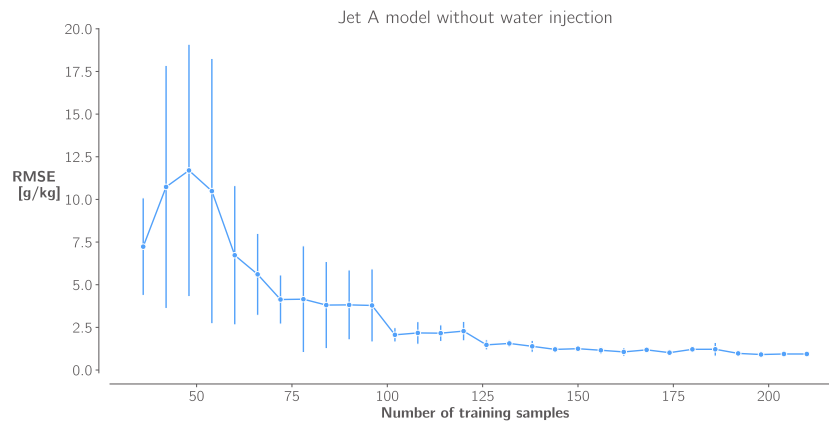


Figure 5.1: The RMS error of the Jet A model without water injection and its deviation starts to flatten at $21 \times n$ sampling points (126 total sampling points).

For this model, the error starts to flatten at a number of 126 sampling points ($21 \times n$), which is well aligned with the recommendations from [76]. Thus, the number of sampling points was set to 126. Since the original sampling dataset used in the cross-validation study was obtained for 210 points ($35 \times n$), a new Latin hypercube sampling with 126 points was generated for the training of the final model, instead of just taking 126 samples out of the 210. This is to ensure that the maximum condition of the Latin hypercube sampling is met in the final set of points, such that the sampling data quality is as high as possible.

Model with water recirculation

For the wet version of the model, a separate surrogate model was created, including the fuel mass flow as an additional variable. The cross-validation results are shown in Figure 5.2. In this case, the error shows a minimum at 125 points before increasing again, so was the number of points chosen for this model.

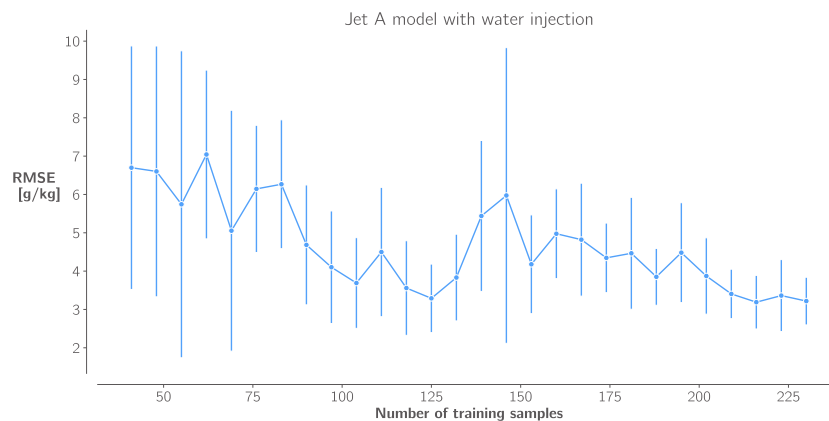


Figure 5.2: The RMS error of the Jet A model with water injection shows a minimum at 125 sampling points before increasing again.

5.4. H₂ Surrogate Model

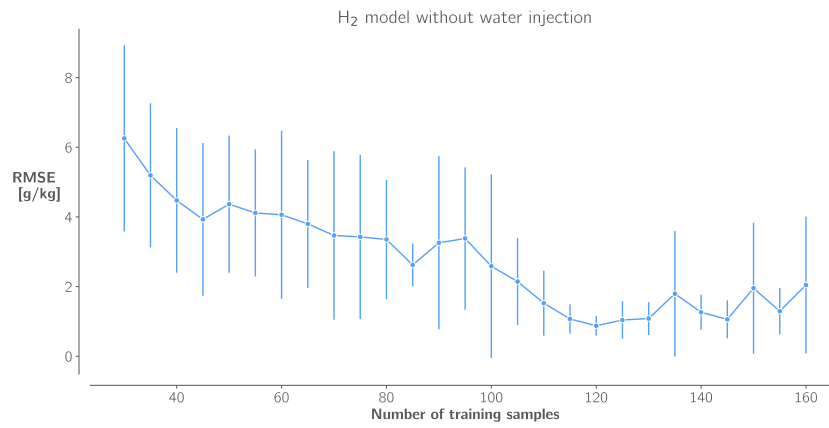
The same procedure was followed for the hydrogen combustor model. The input variables for this model and its limits are shown in Table 5.2. The bounds have been adjusted to the expected range that will be encountered during the optimisation. For example, since hydrogen can burn at much lower equivalence ratios than kerosene, the bounds of the FAR are lower. As described in Chapter 4, the dimensions of this combustor are based on the scaling of the cross-sectional area, and thus the combustor length is no longer a separate input variable.

Table 5.2: Bounds of the sampling intervals for each of the input functions of the H₂ surrogate model.

Parameter	Unit	Lower bound	Upper bound
T_3	K	600	1000
p_3	Pa	14e5	25e5
\dot{m}_{air}	kg/s	8	20
FAR	—	0.005	0.012
A	m ²	0.02	0.08

5.4.1. Cross-validation and Number of Sampling Points

The cross validation study of the H₂ model was performed with the same procedure used for the Jet A models. The results are shown in Figure 5.3. For this model, the error starts to flatten at $24 \times n$ sampling points (120 total points), so this was the chosen number of sampling points.

**Figure 5.3:** The error of the H₂ model without water injection starts to flatten at $24 \times n$ sampling points (120 total sampling points).

5.5. Integration with pyCycle

To integrate the models with pyCycle, each surrogate model was wrapped in an OpenMDAO explicit component, where the `compute` method contains the model output information, and the `compute_partials` method contains the model gradient information, which is available natively in surrogate models created with the SMT.

This component is then added to the pyCycle problem, where the necessary inputs are linked from the relevant pyCycle components, and the output value of the component is used as a constraint in the optimisation problem.

pyCycle Engine Model and Optimisation

Chapter 4 describes the creation of Cantera models to predict engine emissions. Chapter 5 explains the procedure followed to create surrogate models of the Cantera models such that they can be used efficiently in gradient-based optimisation. To test the proposed methodology and models, several optimisations were performed comparing different engine architectures and optimisation constraints. This chapter describes the engine model used in the optimisation problem, as well as the set-up of the optimisation problem itself.

6.1. Baseline Engine

The baseline engine taken as a reference is the N+3 technology level engine by NASA [5]. This engine was developed by NASA as a baseline to evaluate the impact of new engine technologies, making it a good reference to evaluate the impact of new technologies such as water recirculation, or the effect of emissions constraints on the cycle parameters. This engine was also used as the starting point in the work of Atma et al. [27], and as a test case for pyCycle [4], the tool this project is based on.

The engine is a geared, two-spool, high-bypass turbofan engine, and is based on the engine of a single-aisle airliner with the expected turbine engine technology developments for 2025 (N+3 level in the NASA Subsonic Transport System Level Metrics) [5]. The architecture of the engine is shown in Figure 6.1.




Figure 6.1: Architecture of the two-spool turbofan engine.

The architecture above will be compared to an architecture with a water recirculation system, in which water vapour is extracted downstream of the low pressure turbine, and then injected upstream of the high pressure compressor. The implementation and modelling of this system in pyCycle that was done by Atma et al. [27]

will be used in this work.




Figure 6.2: Architecture of the two-spool turbofan engine with the water recirculation system from Atma et al. [27].

6.2. Multipoint Problem Formulation

The optimisation problem set-up is based on the set-up of the test case from the pyCycle paper [4], and the one used in the water recirculation paper from Atma et al. [27], with the addition of EINO_x constraints. The baseline cycle architecture is based on the N+3 engine optimisation problem proposed by Hendricks and Gray [4] to compare pyCycle against NPSS. It is a multi-design point (MDP) approach with four design points. The on-design point is top of climb condition (TOC), and the off-design points are rolling take-off (RTO), cruise (CRZ), and sea level static (SLS). The flight conditions for these points is shown in Table 6.1.

Table 6.1: Flight conditions at each of the four design points.

	Point	Description	Height	Mach number
On-design point	TOC	Top of climb	35000 ft	0.8
Off-design points	RTO	Rolling take-off	0 ft	0.25
	CRZ	Cruise	35000 ft	0.8
	SLS	Sea level static	0 ft	0

The overall set-up of the problem has three levels: the optimisation loop, the MDP cycle, and the individual (off-)design points. The set-up of the problem follows a multidisciplinary feasible (MDF) approach, in which the multidisciplinary analysis (MDA) is converged by a solver for every iteration of the optimiser, such that every iteration of the optimiser results in a feasible design. Note that this feasible design might not meet the optimisation constraints, nor be optimum for every optimiser iteration, but it will satisfy the balance equations and provide a valid design. More background information about this architecture, and a comparison with other MDO architectures, is available in the survey of architectures performed by Martins and Lambe [77].

The balance relationships have the form of residuals, as shown in Equation 6.1, which are converged by a Newton solver inside pyCycle.

$$r(\mathbf{u}) = 0 \quad (6.1)$$

Each of the individual points has its own balance relationships. For the design point, in this case TOC, the balance relationships ensure that the design variables, such as pressure ratios, temperatures, or efficiencies,

are met. This sizes the engine areas, scales the turbomachinery maps, etc. For the off-design points, the balances ensure that fixed engine parameters, such as areas, remain consistent with the sizes calculated in the design point. They also ensure that other prescribed inputs to the off-design point, such as thrust levels, are met too.

The middle loop, the MDP cycle, ensures that the design and off-design points described above are consistent with the overall problem description. The balance relationships in this case prescribe inputs such as thrust ratios between two conditions. More details about the specific balance relationships set up for the N+3 engine cycle, both for the MDP cycle and the individual points, are available in the pyCycle paper [4].

Lastly, the outer loop, the optimisation algorithm, feeds design variables to the MDP cycle, and converges the problem such that the prescribed optimisation constraints are met, and the objective is minimised. The optimisation objectives and constraints used in this project are described in the next section.

6.3. Optimisation Problem

The architecture of the optimisation is shown in the eXtended Design Structure Matrix (XDSM) of Figure 6.3. Vertical lines represent inputs, and horizontal lines represent outputs. The pyCycle MDA component has the multipoint pyCycle problem set-up described in the previous section, and its corresponding XDSMs can be found in [4].

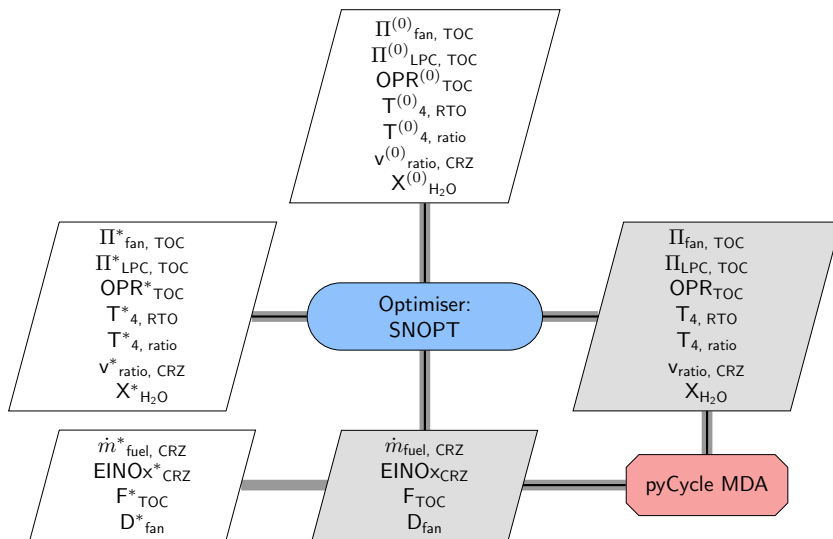


Figure 6.3: Extended design structure matrix (XDSM) of the optimisation problem.

The set-up of the optimisation problem is shown in Table 6.2. The optimisation objective for all conditions is the fuel mass flow during cruise. The optimisation problem has six design variables. The first three of them are the pressure ratios at top of climb of the fan, low-pressure compressor, and the overall one. The fourth design variable is the turbine inlet temperature at the RTO condition, which is the condition of highest turbine inlet temperature. The fifth design variable is the ratio between the turbine inlet temperatures at top of climb and rolling take-off, which dictates the turbine inlet temperature at the TOC condition with respect to the RTO one (which is a design variable itself). This ratio is shown in Equation 6.2.

$$T_{4, \text{ratio}} = \frac{T_{4, \text{TOC}}}{T_{4, \text{RTO}}} \quad (6.2)$$

The sixth design variable is the nozzle speed ratio during cruise, which is defined by Equation 6.3. This design variable effectively controls the bypass ratio of the engine. The C_v coefficients are the nozzle velocity

coefficients that correct the ideal nozzle speeds for non-ideal effects, and are fixed inputs.

$$v_{\text{ratio}} = \frac{v_{\text{core, ideal}} C_{\nu, \text{core}}}{v_{\text{bypass, ideal}} C_{\nu, \text{bypass}}} \quad (6.3)$$

The last design variable is the mass fraction of water extracted from the flow (e.g., a 0.5 fraction means that 50% of the water vapour present in the exhaust flow is extracted). Depending on the specific case, as further detailed below, this variable might represent only the extraction fraction in cruise, or across the four design points. However, it will be a single design variable in all cases. This design variable is of course only present in the cases with water recirculation.

Table 6.2: Description of the optimisation problem.

	Variable	Description	Unit	Quantity
minimise	$\dot{m}_{\text{fuel, CRZ}}$	Fuel mass flow during cruise	kg/s	1
by varying	$\Pi_{\text{fan, TOC}}$	Fan pressure ratio at top of climb	—	1
	$\Pi_{\text{LPC, TOC}}$	LPC pressure ratio at top of climb	—	1
	OPR_{TOC}	Overall pressure ratio at top of climb	—	1
	$T_{4, \text{RTO}}$	Turbine inlet temperature at rolling take-off	K	1
	$T_{4, \text{ratio}}$	Ratio of T_4 at TOP and RTO	—	1
	$v_{\text{ratio, CRZ}}$	Ratio of nozzle speeds in the core and bypass	—	1
	$X_{\text{H}_2\text{O}}$	Water recovery mass fraction	—	1
		Total		7
subject to	$F_{\text{TOC}} \geq 25.8$	Thrust at top of climb	kN	1
	$D_{\text{fan}} \leq 2.54$	Fan diameter	m	1
	$\text{EINO}_{\text{XCRZ}} \leq 60\% \text{ EINO}_{\text{XCRZ}, 0}$	NO_x emissions index during cruise	g/kg	1
		Total		3

The design variables are allowed to change within the bounds described in Table 6.3.

The optimisation is constrained with three inequality constraints. The first one is that the thrust level at top of climb must be at least 25.8 kN, to ensure that the optimised engine still meets performance requirements. The second constraint limits the fan diameter to a maximum of 2.54 m (100 inch), which is the size of reference N+3 engine [5]. The third and last constraint is the NO_x emissions index (EINO_x) during cruise. This constraint varies in magnitude for each of the optimisation cases. First, each case is optimised for fuel consumption without any EINO_x constraint, and then each case is re-optimised limiting the EINO_x levels to 60% of the value obtained in the unconstrained optimisation. This is done such that the effects that limiting emissions has on engine performance can be compared. The reason why a relative constraint is used, and not an absolute one, is that the order of magnitude of the EINO_x of Jet A engines is around $20\times$ higher than for H_2 engines; thus, a constraint that is below the unconstrained EINO_x of H_2 would be unachievable by the Jet A engine.

6.3.1. Optimisation Cases

The optimisation architecture described until now describes the general approach applicable to all cases. However, some small fundamental differences in the engine architectures mean that two different optimisation approaches are chosen for the dry and wet engines. This section contains a list of all the cases included in the results section.

Table 6.3: Bounds of the design variables in the optimisation problem.

Variable	Unit	Lower bound	Upper bound
$PR_{fan, TOC}$	—	1.2	1.4
$PR_{LPC, TOC}$	—	2.5	4.0
OPR_{TOC}	—	40.0	70.0
$T_{4, RTO}$	K	1667	2000
$T_{4, ratio}$	—	0.5	0.95
$V_{ratio, CRZ}$	—	1.35	1.45
X_{H_2O}	—	0	0.7

For the dry engines, they are first optimised for minimum fuel flow during cruise, with no NO_x constraint. The NO_x emissions at the optimised value are taken as a reference, and the engines are optimised again with an added NO_x constraint, set to 60% of the baseline value for each of the two engines. This is carried out separately for two different fuels, Jet A and H₂. This corresponds to the first four cases in Table 6.4.

For the wet engine, a different approach was chosen. The objective of the cases with water recirculation was to understand whether the optimiser would attempt to reduce emissions by modifying the engine design through the first six design variables (as in the dry cases), or if it would rely mostly on adding extra water. For this, the engine is first optimised with the same set up as the dry engine, but with water extraction fraction X_{H_2O} as an added design variable. In this first case, the fraction was set to be equal across all four design points (which implies different water mass flows in each point). This case is referred to as case *JW*. This set-up is not necessarily optimal and is expected to yield a worse performing engine than if each of the extraction fraction was optimised separately. However, it was chosen because 1) it provides a good baseline for the second case, explained below; 2) since only the fuel consumption or NO_x emissions of the cruise point are part of the objective and constraints, the benefit of adding water in other off-design points was not properly captured by the optimiser. In a more realistic and complex case, NO_x emissions would also be constrained in other flight phases, such as take-off, and the fuel consumption in other phases would probably be included in a multi-objective optimisation approach. Thus, forcing the extraction fraction to be the same across all four points provides a way to ensure water flow across all design points.

The NO_x emissions of the optimised configuration from this case are taken as a reference, and the engine is again re-optimised with an added NO_x constraint set to 60% of the baseline value from case *JW*. However, the design variables are modified in the re-optimisation. The second optimisation is performed fixing the extraction fraction X_{H_2O} in TOC, RTO, and SLS to the value found in case *JW*, and only the cruise extraction fraction is optimised. The reasoning behind this is to focus on the effect that adding water has on the performance and emissions of an off-design point (CRZ), while limiting the effect that the water injection has on the sizing of the engine in the design point (TOC). Moreover, since the extraction fraction in cruise is allowed to vary independently from the fraction at other design points, this allows to potentially have a case with both lower emissions and better performance than case *JW*, were the fraction was the same across all points. This case is referred to as case *JWNa*.

Then, a third optimisation is performed; in this case, the extraction fraction across all four points is fixed to the value found in case *JW*, and the engine is optimised with only the six design variables used in the dry cases, but with the presence of a fixed amount of water recirculation and with the NO_x constraint set to 60% of the EINO_x in case *JW*. This case is referred to as case *JWnb*. This case was added mostly to show the benefit that extra water extraction (case *JWNa*) has with respect to simply modifying the engine architecture to reduce emissions. The expected result is to have the same emissions as in case *JWNa* but with a higher fuel flow.

Table 6.4: List of the different optimisation cases of this project, along with their associated codes.

Case code	Fuel	NOx constraint	Water fraction as design variable
<i>JD</i>	Jet A	None	No water extraction
<i>JDN</i>	Jet A	60% of <i>JD</i>	No water extraction
<i>HD</i>	H ₂	None	No water extraction
<i>HDN</i>	H ₂	60% of <i>HD</i>	No water extraction
<i>JW</i>	Jet A	None	Yes – X _{H₂O}
<i>JWNa</i>	Jet A	60% of <i>JW</i>	Yes – X _{H₂O} , CRZ
<i>JWNb</i>	Jet A	60% of <i>JW</i>	No – X _{H₂O} fixed to value of case <i>JW</i>

6.4. Optimiser

The optimiser used in this project is SNOPT (Sparse Nonlinear OPTimizer) [6]. SNOPT is an SQP (sequential quadratic programming) algorithm for constrained, gradient-based optimisation of linear and nonlinear problems, and is well-suited for large, sparse problems. SNOPT is used in pyCycle through the pyOptSparse interface [78], which is an open-source Python package that facilitates constrained, nonlinear optimisation, and provides a unified interface for several optimisation algorithms, including SNOPT. pyOptSparse is integrated into OpenMDAO (which pyCycle is based on), thus allowing direct use of SNOPT in pyCycle.

Table 6.5: Optimality and feasibility tolerances for all optimisation cases.

	Optimality	Feasibility
<i>Tolerance</i>	10 ⁻⁶	10 ⁻⁶

Optimisation Results and Comparison

This chapters contains the results of the optimisation cases presented in Chapter 6. The results will be presented in four parts. The first part shows the optimisation results of the baseline engine architecture with both fuels (Jet A and H₂). The second part contains the optimisation results of the engine with water recirculation. The third section contains a comparison of the engine performance of all seven cases. The fourth and last section shows the computational aspect of the simulations, specifically the simulation time and number of iterations, and the effects that adding NO_x constraints has on these.

7.1. Results of the Dry Cases

The results of the cases of the conventional turbofan engine without water recirculation are shown in Table 7.1. As explained in Chapter 6, the optimisation was first run for each engine without a NO_x constraint (the columns labelled as 'Baseline'). The NO_x emissions in those cases where then calculated, and the optimisations were then re-run adding a constraint limiting the EINO_x to be 60% of the baseline value (the columns labelled as '60% NO_x'). This approach of limiting the EINO_x relative to the baseline value was chosen over a common constraint for all engines because the NO_x emissions of the engines with Jet A and H₂ have different orders of magnitude.

Table 7.1: Results of the optimisations of the cases without water recirculation. A lower asterisk * represents an active lower bound or \geq constraint, and a high asterisk represents an active upper bound or \leq constraint. The values in square brackets are reported for completeness, but were not implemented as constraints.

		Jet A - Dry		H ₂ - Dry		
		Baseline	60% EINO _x	Baseline	60% EINO _x	
		(JD)	(JDN)	(HD)	(HDN)	
<i>Objective</i>	$\dot{m}_{\text{fuel, CRZ}}$	kg/s	0.2902	0.2936	0.10580	0.10584
<i>Constraints</i>	EINO _{x,CRZ}	g/kg	[25.12]	15.07*	[1.214]	0.729*
	F _{TOC}	kN	25.8*	25.8*	25.8*	25.8*
	D _{fan}	m	2.54*	2.54*	2.54*	2.54*
<i>Design variables</i>	$\Pi_{\text{fan, TOC}}$	—	1.285	1.288	1.2883	1.2866
	$\Pi_{\text{LPC, TOC}}$	—	4*	4*	4*	4*
	OPR _{TOC}	—	60.71	44.29	60.385	61.706
	T _{4, RTO}	K	1796.2	1807.1	1786.7	1737.0
	T _{4, ratio}	—	0.9102	0.9106	0.9154	0.9141
	V _{ratio, CRZ}	—	1.35*	1.35*	1.35*	1.35*

For the Jet A case, reducing the NO_x emissions by 40% only resulted in a 1.18% increase in fuel mass flow.

This is an important result, because it shows that including emissions constraints during preliminary design can lead to engine designs with reduced emissions without noticeable performance penalties.

For the H₂ case, a 40% reduction in NO_x emissions is much smaller in absolute value than for the Jet A case, since the baseline emissions are lower. Thus, the increase in NO_x emissions only represents a 0.036% increase in fuel consumption, which is marginal.

It is interesting to note that all constraints are active in all four cases, showing that the optimiser always tends to increase the fan diameter, and to reduce the thrust in the design point to the minimum. A few of the design variable bounds are also always active, which is the case for the LPC pressure ratio (active upper bound), and the nozzle velocity ratio (active lower bound).

The design variables, however, correspond to parameters of the design point (TOC) and RTO. The NO_x emission constraint applies only to cruise, so these design variables by themselves do not provide intuition into what is causing the reduced emissions in cruise condition. Table 7.2 shows the burner conditions during cruise, which are the inputs used by the surrogate model used to predict NO_x emissions.

Table 7.2: Inputs to the emissions surrogate models for each of the cases without water recirculation.

Variable	Unit	Jet A - Dry		H ₂ - Dry	
		Baseline (JD)	60% EINO _x (JDN)	Baseline (HD)	60% EINO _x (HDN)
T ₃	K	852.73	772.54	855.67	858.32
p ₃	Pa	20.19 · 10 ⁵	14.72 · 10 ⁵	20.09 · 10 ⁵	20.54 · 10 ⁵
\dot{m}_{air}	kg/s	13.86	12.57	12.89	13.95
FAR	—	0.0209	0.0234	0.00821	0.00759
ϕ_{PZ}	—	1.0737	1.2021	—	—
ϕ_{SZ}	—	0.5497	0.6155	—	—
ϕ	—	0.3010	0.3370	0.2791	0.2581
A	m ²	0.0437	0.0547	0.0414	0.0430
L	m	0.1835	0.1751	—	—
EINO _x	g/kg	25.12	15.07	1.214	0.729

Looking at the results for the Jet A case, the most significant change is the reduction of the inlet temperature (T₃). This reduces the maximum temperature reached in the combustor, and thus reduces thermal NO_x. The fuel-to-air ratio therefore increases to compensate, but the overall result produces less NO_x. The combustor inlet pressure is reduced by around 25%. The reason for the reduction in pressure is mainly because compression also implies a raise in temperature, so the optimiser reduces T₃ by lowering the overall compression ratio, rather than the direct effect of pressure on NO_x formation, which is weak.

For the H₂, since the absolute value of emissions does not need to reduce by, the overall design of the engine stays reasonably similar between both optimisations. The reduction in NO_x is likely due to the reduced FAR in cruise, which causes a reduction in burner exit temperature (given that the inlet temperature only differs by 3 K between both cases).

7.2. Results of the Cases with Water Recirculation

The next set of results corresponds to the turbofan engine with water recirculation. First, the engine was optimised using the same design variables as in the dry cases, with the addition of the water extraction fraction (setting the extraction fraction in all four design points to be the same), and the EINO_x was calculated at this point (case *JW*).

In the second case (case *JWNa*), the optimisation was repeated but varying only the water extraction fraction in cruise (along with the other design variables used in the dry cases), while fixing the water extraction fraction for TOC, RTO, and SLS to the optimised value found in case *JW*. For this case, the NO_x emissions were constrained to be at most 60% of those in case *JW*. The idea behind this approach is to see whether the optimiser reduces the emissions by varying the engine design mainly through the original design variables (following the same trends as for the dry engine), or if it instead relies mainly on injecting more water.

The third case (case *JWNB*) corresponds to an optimisation using only the design variables from the dry cases, while fixing all four water extraction ratios to the value found in case *JW*. As for the previous case, the NO_x emissions were constrained to be at most 60% of those in case *JW*. This case is mainly intended as a comparison case, to see the difference in performance when the optimiser is not allowed to use more water to reduce the emissions. These results are shown in Table 7.3.

Table 7.3: Results of the optimisations of the cases with water recirculation. A lower asterisk * represents an active lower bound or \geq constraint, and a high asterisk represents an active upper bound or \leq constraint. The values in square brackets are reported for completeness, but were not implemented as constraints or design variables.

		Jet A - Wet			
			<i>Baseline</i>	60% <i>EINO_x</i>	60% <i>EINO_x</i>
			<i>(JW)</i>	<i>(JWNa)</i>	<i>(JWNB)</i>
<i>Objective</i>	$\dot{m}_{\text{fuel, CRZ}}$	kg/s	0.2896	0.2874	0.2912
<i>Constraints</i>	$EINO_{\text{XCRZ}}$	g/kg	[31.58]	18.95*	18.95*
	F_{TOC}	kN	25.8*	25.8*	25.8*
	D_{fan}	m	2.54*	2.54*	2.54*
<i>Design variables</i>	$\Pi_{\text{fan, TOC}}$	—	1.283	1.288	1.287
	$\Pi_{\text{LPC, TOC}}$	—	4*	4*	4*
	OPR_{TOC}	—	56.77	40.00*	45.44
	$T_{4, \text{RTO}}$	K	1796.86	1849.23	1870.38
	$T_{4, \text{ratio}}$	—	0.9049	0.9052	0.9065
	$V_{\text{ratio, CRZ}}$	—	1.35*	1.35*	1.35*
	$X_{\text{H}_2\text{O, CRZ}}$	—	0.1027	0.1723	[0.1027]
	$X_{\text{H}_2\text{O, TOC/RTO/SLS}}$	—	0.1027	[0.1027]	[0.1027]

The trends of the design variables is similar to that observed in the dry cases, with all constraints being active. The cases with constrained NO_x emissions show lower overall pressure ratios, and higher T_4 at the RTO point. For case *JW*, where the extraction ratio is set to be equal across all four design points, the optimum extraction ratio is 0.1027 (i.e., 10.27% of the water vapour present in the exhaust gas is condenses and recirculated). This value is rather low, but this result is caused by the different extraction ratios not being allowed to vary independently, added to the fact that the objective and the NO_x constraint only correspond to one off-design point. For the second case, where the cruise extraction fraction is allowed to vary independently, this value increases up to 0.1723. This results not only in the lower emissions, but also in a lower fuel consumption during cruise, which means that the engine has both better performance and better emissions. Of course, this does not mean that the optimum value for case *JW* was not a real optimum, but rather it implies that allowing the water extraction ratios to vary independently leads to better performance and lower emissions. It also shows that water extraction can be a way to reduce emissions while also increasing the performance of the engine.

As for the dry cases, the cause of the reduction in NO_x emissions during cruise can be better understood by analysing the combustor conditions during cruise, which are shown in Table 7.4.

Looking at Table 7.4, both cases with NO_x constraints (*JWNa* and *JWNB*) have a noticeably lower combustor inlet temperature than the case optimised without NO_x constraints, which has an important effect on reducing

Table 7.4: Inputs to the emissions surrogate models for each of the cases with water recirculation.

Variable	Unit	Jet A - Wet		
		Baseline (<i>JW</i>)	60% EINO _x (<i>JWNa</i>)	60% EINO _x (<i>JWnb</i>)
T ₃	K	889.99	851.45	842.58
p ₃	Pa	18.89 · 10 ⁵	13.27 · 10 ⁵	15.10 · 10 ⁵
\dot{m}_{air}	kg/s	14.69	12.50	12.59
FAR	—	0.0197	0.0230	0.0231
ϕ_{PZ}	—	1.0121	1.1816	1.1867
ϕ_{SZ}	—	0.5181	0.6049	0.6076
ϕ	—	0.2837	0.3312	0.3326
A	m ²	0.0493	0.0614	0.0543
L	m	0.1888	0.1808	0.1839
\dot{m}_{water}	kg/s	0.0366	0.0659	0.0361
EINO _x	g/kg	31.58	18.95	18.95

emissions. The case where the water extraction in cruise is allowed to vary, the water mass flow is close to double, which also contributes to lowering emissions (given that the overall mass flow is smaller, and the FAR increases in a lower proportion). For the case where the water extraction is fixed, the combustor inlet temperature is the lowest, to achieve the same emissions as in case *JWNa* without the extract water flow.

7.3. Comparison of Cases

The optimisation objective of both the hydrogen and Jet A cases was the fuel flow during cruise. However, the magnitude of this objective is approximately three times bigger for Jet A than for H₂ due to their different energy densities. This also implies that their thrust-specific fuel consumption (TSFC) values will differ in magnitude. Thus, to do a more accurate analysis of the results between the two fuels, a different metric will be used: the thrust-specific energy consumption. This metric, proposed by Adler and Martins [8], combines the conventional TSFC [kg/N/s] with the lower heating value (LHV) [J/kg] of the fuel, resulting in the thrust-specific energy consumption (TSEC) [J/N/s] shown in Equation 7.1. The LHV of Jet A is 43 MJ/kg, and the LHV of H₂ is 120 MJ/kg.

$$\text{TSEC} = \text{TSFC} \cdot \text{LHV} \quad (7.1)$$

The TSFC and TSEC of each case is shown in Table 7.5. As it can be seen from the results, and as was shown in previous sections, implementing NO_x constraints increases the TSEC as expected. This happens in all cases except one, case *JWNa*. In this case, the addition of extra water extraction during cruise yields a better performance on top of the reduced emissions, which shows the benefits of water recirculation. The results also show that the cases with water recirculation have lower TSEC and TSFC than the dry cases. However, the cases were not set up to evaluate the full potential of water recirculation, but to evaluate the NO_x constraints on an engine with water recirculation. Thus, it is expected that if the water extraction fractions are allowed to vary independently, even lower TSEC can be achieved in the water recirculation cases, and as a result these TSEC values should not be used as an absolute comparison between dry and wet engines.

Figure 7.1 shows the percentual difference in TSEC between each case with a NO_x constraint compared to its corresponding baseline case. As expected, constraining NO_x emissions leads to a lower engine efficiency, except for case *JWNa*, where the increase in the water extraction fraction during cruise results in a lower TSEC in addition to the reduced emissions. The increase in TSEC in the hydrogen case is much lower

Table 7.5: Thrust-specific fuel and energy consumption of each of the seven cases analysed.

		Jet A - Dry		H ₂ - Dry	
		Baseline (JD)	60% EINO _x (JDN)	Baseline (HD)	60% EINO _x (HDN)
TSFC	[kg/N/s]	$1.2498 \cdot 10^{-5}$	$1.2645 \cdot 10^{-5}$	$4.5565 \cdot 10^{-6}$	$4.5582 \cdot 10^{-6}$
TSEC	[J/N/s]	537.40	543.75	546.78	546.98

		Jet A - Wet		
		Baseline (JW)	60% EINO _x (JWNa)	60% EINO _x (JWNa)
TSFC	[kg/N/s]	$1.2470 \cdot 10^{-5}$	$1.2377 \cdot 10^{-5}$	$1.2543 \cdot 10^{-5}$
TSEC	[J/N/s]	536.22	532.21	539.33

than in the Jet A case, but this a result of the fact that the hydrogen engine has a much lower baseline emissions index than the Jet A case, so limiting the NO_x constraint to 60% of the baseline value implies a low reduction in NO_x in absolute terms, which results in a very similar engine design in both the constrained and unconstrained cases. Case *JWNa*, where the water extraction fraction is kept constant, results in half the increase in TSEC compared to the dry case with Jet-A.




Figure 7.1: This figure shows the relative change in TSEC of each case with a NO_x constraint compared to its corresponding baseline case. Implementing NO_x constraints results in higher specific fuel and energy consumption except in case *JWNa*, where the addition of water during cruise offers an increase in efficiency while reducing emissions.

7.4. Computational Time and Number of Iterations

To close the results, it is important to analyse the effect that adding the NO_x models has in the computation, in terms of number of iterations, and optimisation time.

Figure 7.2 shows the evolution of the optimality and the feasibility of the four dry cases for each major iteration. Adding the NO_x constraint increases the number of major iterations for both the Jet A and H₂ cases. However, when looking at the computational time, shown in Table 7.6, adding the NO_x constraint only increases the computational time by 3 minutes in the Jet A case, and results in a lower computational time for the H₂ case.




Figure 7.2: The addition of the surrogate model and the associated NO_x constraint results in more iterations for both dry cases.

For the cases with water recirculation, the evolution of the optimality and feasibility is shown in Figure 7.3. The wall time for these cases is also shown in Table 7.6. Each of these cases has different design variables, and different starting points (cases JWNa and JWNb are started from the optimum point found in case JW, since the water extraction fraction in JWNa and JWNb is set to that value), meaning that the comparison of number of iterations and computational time between these cases is not as insightful as for the dry cases.




Figure 7.3: The addition of the surrogate model and the associated NO_x constraint results in less iterations for the case in which the water extraction fraction is fixed (JWNb), but more iterations for the case in which the cruise extraction fraction is a design variable (JWNa).

The reported time is the wall time of running the optimisations in a standard laptop with 6 Intel i7 cores, and 16 GB of RAM. This wall time is only orientative, since a proper benchmark was not carried out, and the CPU load while running each case was not collected. However, it still shows that adding a surrogate model and a NO_x constraint does not have a noticeable, negative effect on the computational cost of the optimisations, with 4 out of the 6 optimisations using NO_x constraints have lower wall times than the corresponding optimisations without a NO_x constraint.

Table 7.6: Wall time of the optimisations, run in a standard laptop with 6 cores and 16 GB of RAM.

		Jet A - Dry		H₂ - Dry	
		<i>Baseline</i>	<i>60% EINO_x</i>	<i>Baseline</i>	<i>60% EINO_x</i>
Wall time	[min]	35	39	41	28

		Jet A - Wet		
		<i>JW</i>	<i>JWNa</i>	<i>JWNb</i>
Wall time	[min]	59	55	51



Conclusion

Current turbofan engine design trends are approaching practical limits due to maximum allowable material temperatures in the turbines, and due to the exponential increase in NO_x emissions at very high temperatures. To solve this, new engine architectures or new fuels need to be considered, and NO_x emissions need to be part of the engine design cycle starting from its preliminary design.

To reach this goal, the research objective of this thesis was to

investigate the effects of including NO_x emission constraints in gradient-based optimisation of conventional and novel turbofan engines.

To achieve this, it was decided to use simplified chemical reactor networks to predict emissions. The use of simplified network architectures allows to use them during preliminary design without the need of detailed combustor or flow field geometry, posing a big advantage over CFD simulations. At the same time, these networks still allow for great flexibility in terms of inputs and operating conditions, so as to effectively analyse both conventional and unconventional designs with a level of accuracy that would not be possible with simpler approaches such as analytical correlations.

Two separate models were created in this work, one for conventional Jet A combustors, based on an architecture proposed by Villette et al. [1] consisting of three perfectly stirred reactors corresponding to the primary, secondary, and dilution zones of the reactor; and one for hydrogen lean premixed combustors, based on an architecture proposed by Talboom [2] consisting of four perfectly stirred reactors and one plug-flow reactor.

To use these models in a gradient-based optimisation framework, a surrogate model was created for each of them. The surrogate models were created using the Surrogate Modelling Toolbox [3], which provides analytical gradient information of the models. The models used the Kriging method, and the sampling dataset was created with Latin hypercube sampling.

To test the approach in an optimisation, a multipoint optimisation of a turbofan engine was carried out with pyCycle [4]. The baseline engine is based on the N+3 technology reference level by NASA [5], which consists of a geared, two-spool, high bypass turbofan. The engine was first optimised separately for both Jet A and hydrogen as fuels, to test the proposed methodology with conventional and unconventional fuels. Additionally, an engine with a closed-loop water recirculation system was optimised to test the approach with an unconventional architecture. For all of these cases, the optimisation was first performed without any emissions constraints, and then a second optimisation was carried out limiting the NO_x emissions to 60% of the baseline value for each case.

The multipoint optimisation problem consisted of one design point (top of climb), and three off-design points (take-off, sea level static, and cruise). The optimisation was performed using SNOPT [6], an SQP algorithm for large, constrained, gradient-based optimisation problems.

The takeaways of this research were the following:

- Simplified chemical reactor networks provide a good balance between accuracy and simplicity, allowing more freedom during preliminary design than empirical correlations. At the same time, they do not rely on detailed combustor geometry or flow field information that is necessary for more complex reactor network approaches. (See Chapter 3 and Chapter 4.)
- Surrogate modelling allows to integrate chemical reactor networks with gradient-based tools to overcome the lack of analytical gradients in tools like Cantera. They also offer other benefits, such as faster model probing, and avoid convergence issues with the reactor networks. (See Chapter 5.)
- The optimisations achieved 40% reduction in NO_x with a maximum penalty in fuel consumption of only 1.2%. This means that significant reductions in NO_x can be achieved with minimal penalties in fuel consumption, which supports the use of NO_x analysis and constraints during preliminary engine design. (See Chapter 7.)
- In the analysed engine configurations, H₂ leads to NO_x emissions that are 1 or 2 orders of magnitude below those of Jet A engines. (See Chapter 7.)
- Water recirculation can reduce NO_x emissions at the same time as improving engine performance. (See Chapter 7.)

This thesis provides a working framework for gradient-based optimisation of conventional and novel engine architectures including NO_x emissions constraints. The proposed emissions models can be used as-is for optimisation of engines using conventional Jet A combustors, or lean, premixed hydrogen combustors. The presented test cases can be used as reference points to test engine architecture modifications.

Furthermore, the proposed framework offers ample flexibility at all levels. It allows independent modification of the reactor networks, the surrogate modelling approach, or the optimisation problem. For example, the proposed reactor networks can be used to optimise new engine configurations, or entirely new reactor networks based on different combustor types can be tested using the proposed optimisation set-up.

While this project was focused on preliminary design, where limited information about combustor geometry is usually available, the methodology can be used with arbitrarily more complex CRNs whenever the required information is available to the user.

Future Recommendations

Developing methods and frameworks for preliminary design inherently requires a trade-off between applicability and accuracy. The methods presented here are intended to offer the needed freedom during preliminary design while maintaining as much accuracy in the trends as possible. However, this still comes with some limitations in the framework. To conclude the work presented in this project, some recommendations for improvements and future work are presented here.

- The chemical reactor network for the conventional Jet A combustor has four combustor-dependent parameters that are expected to change for each engine. Since these parameters are not available from pyCycle, they were tuned using the data from the three engines analysed in the original work where the architecture was presented [1]. However, since this tuning is meant to generalise the applicability of the network to any arbitrary engine, extending the tuning dataset from three to more engines would potentially increase the accuracy of the NO_x trends of the model.
- The chemical reactor network used for the hydrogen cases was based on a network that was built and tuned to match the results of a experiment that was carried out in operating conditions different from those of a gas turbine combustor (e.g., it was carried out at ambient pressure). While some adjustments were made and the network was scaled to match the size of the combustor, the results of this network could not be compared against real data from a full-sized combustor under normal operating conditions. This kind of data for hydrogen combustors is currently very scarce, at least from publicly available sources, but when more experimental data becomes available in the future, further tuning and validation of the reactor network would likely increase the accuracy of the trends and the confidence in the results.
- The water recirculation system used in this work was modelled in pyCycle by another author in a previous project [27]. As explained in the original work, this model is a simplified implementation of a water condenser and an injector. The change of phase of the water along this process and the pressure drop associated with the system are not modelled. A more accurate implementation that calculates the pressure drop as a function of the amount of water extracted would give more insight into the overall effect on the engine of water recirculation.
- The current optimisation set-up only includes cruise fuel consumption as an objective, and only constraints NO_x emissions during the cruise phase. However, a more realistic scenario would include NO_x emission limits also in the take-off phase, or for the whole LTO (landing and take-off) cycle. The goal of this project was focused on the methodology of implementing NO_x constraints in gradient-based optimisation, which was achieved with the proposed cases. However, this more complex set-up would provide more insight into the optimal water extraction fractions across the different flights phases.

References

- [1] Sergios Villette et al. "Simplified Chemical Reactor Network Approach for Aeroengine Combustion Chamber Modeling and Preliminary Design". In: *Aerospace* 11.1 (Jan. 2024), p. 22. DOI: 10.3390/aerospace11010022.
- [2] Mathijs G. Talboom. "Chemical Kinetics Study of the Hybrid Combustion System". MA thesis. Delft University of Technology, Nov. 2016. URL: <http://resolver.tudelft.nl/uuid:bd91a4da-22c5-4ff3-95e4-7a8088c6931c>.
- [3] Paul Saves et al. "SMT 2.0: A Surrogate Modeling Toolbox with a focus on Hierarchical and Mixed Variables Gaussian Processes". In: *Advances in Engineering Software* 188 (2024), p. 103571. DOI: 10.1016/j.advengsoft.2023.103571.
- [4] Eric S. Hendricks and Justin S. Gray. "pyCycle: A Tool for Efficient Optimization of Gas Turbine Engine Cycles". In: *Aerospace* 6.87 (Aug. 2019). DOI: 10.3390/aerospace6080087.
- [5] Scott M. Jones, William J. Haller, and Michael T. Tong. *An N+3 Technology Level Reference Propulsion System*. Tech. rep. NASA/TM—2017-219501. NASA Glenn Research Center, 2017. URL: <https://ntrs.nasa.gov/citations/20170005426>.
- [6] Philip E. Gill, Walter Murray, and Michael A. Saunders. "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization". In: *SIAM Review* 47.1 (2005), pp. 99–131. DOI: 10.1137/S0036144504446096.
- [7] Hannah Ritchie, Pablo Rosado, and Max Roser. "CO₂ and Greenhouse Gas Emissions". In: *Our World in Data* (2023). Accessed on June 26th 2023. URL: <https://ourworldindata.org/co2-and-greenhouse-gas-emissions>.
- [8] Eytan J. Adler and Joaquim R. R. A. Martins. "Hydrogen-Powered Aircraft: Fundamental Concepts, Key Technologies, and Environmental Impacts". In: *Progress in Aerospace Sciences* 141 (Aug. 2023), p. 100922. DOI: 10.1016/j.paerosci.2023.100922.
- [9] Feijia Yin and Arvind Gangoli Rao. "Performance of multi-fuel hybrid engine for blended wing body". In: *Proceedings of the 4th International Symposium on Jet Propulsion and Power Engineering*. Xi'an, China, Sept. 2012, pp. 581–588.
- [10] David L. Daggett et al. *Water Injection on Commercial Aircraft to Reduce Airport Nitrogen Oxides*. Tech. rep. NASA/TM-2010-213179. NASA Glenn Research Center, Mar. 2010.
- [11] Salvatore Iavarone et al. "A multiscale combustion model formulation for NO_x predictions in hydrogen enriched jet flames". In: *International Journal of Hydrogen Energy* 44.41 (Aug. 2019), pp. 23436–23457. ISSN: 0360-3199. DOI: 10.1016/j.ijhydene.2019.07.019.
- [12] Stephen R. Turns. *An Introduction to Combustion: Concepts and Applications*. 2nd ed. McGraw-Hill, 2000. ISBN: 9780071169103.
- [13] European Commission and Directorate-General for Climate Action. *Going climate-neutral by 2050 – A strategic long-term vision for a prosperous, modern, competitive and climate-neutral EU economy*. Publications Office, 2019. DOI: 10.2834/02074.
- [14] European Commission et al. *Flightpath 2050 – Europe's vision for aviation – Maintaining global leadership and serving society's needs*. Publications Office, 2011. DOI: 10.2777/50266.
- [15] United Nations. *Paris Agreement to the United Nations Framework Convention on Climate Change*. Article 2.a). T.I.A.S. No. 16-1104. Dec. 2015.
- [16] Hannah Ritchie. "Cars, planes, trains: where do CO₂ emissions from transport come from?" In: *Our World in Data* (2020). Accessed on June 26th 2023. URL: <https://ourworldindata.org/co2-emissions-from-transport>.
- [17] Hannah Ritchie. "What share of global CO₂ emissions come from aviation?" In: *Our World in Data* (2024). Accessed on June 26th 2023. URL: <https://ourworldindata.org/global-aviation-emissions>.
- [18] Andreas W. Schäfer et al. "Technological, economic and environmental prospects of all-electric aircraft". In: *Nature Energy* 4.2 (Feb. 2019), pp. 160–166. DOI: 10.1038/s41560-018-0294-x.
- [19] Joseph Brand et al. "Potential Use of Hydrogen In Air Propulsion". In: *AIAA International Air and Space Symposium and Exposition: The Next 100 Years*. Dayton, Ohio, July 2003. DOI: 10.2514/6.2003-2879.
- [20] C. Koroneos et al. "Advantages of the use of hydrogen fuel as compared to kerosene". In: *Resources, Conservation and Recycling* 44.2 (2005), pp. 99–113. ISSN: 0921-3449. DOI: 10.1016/j.resconrec.2004.09.004.
- [21] Miguel Ángel Sáez Ortuño et al. "Climate Assessment of Hydrogen Combustion Aircraft: Towards a Green Aviation Sector". In: *AIAA SciTech 2023 Forum*. Vol. AIAA 2023-2513. National Harbor, MD, Jan. 2023. DOI: 10.2514/6.2023-2513.
- [22] Loek K. H. Gijzen. "Exploratory Analysis of Future LH₂-Powered Aircraft Ground Operations at a Regional Airport". MA thesis. Delft University of Technology, Apr. 2024. URL: <http://resolver.tudelft.nl/uuid:1f76e80b-2111-46fb-8eaf-c08e024bd3e7>.
- [23] Jonas Mangold et al. "Refueling of LH₂ Aircraft — Assessment of Turnaround Procedures and Aircraft Design Implication". In: *Energies* 15.7 (2022), p. 2475. ISSN: 1996-1073. DOI: 10.3390/en15072475.
- [24] Luuk ten Damme. "Simulating the Refuelling Process for a Liquid Hydrogen-Powered Commercial Aircraft". MA thesis. Delft University of Technology, Sept. 2023. URL: <http://resolver.tudelft.nl/uuid:6a86c67f-b509-45d5-8b8b-fe6f8802aca7>.
- [25] Feijia Yin. "Modeling and Characteristics of a Novel Multi-fuel Hybrid Engine for Future Aircraft". PhD thesis. Delft University of Technology, 2016. DOI: 10.4233/uuid:344b7d9c-f54c-4836-87ca-28582231a3d3.
- [26] Feijia Yin and Arvind Gangoli Rao. "Performance analysis of an aero engine with inter-stage turbine burner". In: *The Aeronautical Journal* 121.1245 (Sept. 2017), pp. 1605–1626. DOI: 10.1017/aer.2017.93.
- [27] Peter N. Atma, Andrew H. R. Lamkin, and Joaquim R. R. A. Martins. "Comparing Hydrogen and Jet-A for an N+3 Turbofan with Water Recirculation using Gradient-Based Optimizaiton". In: *AIAA Aviation Forum*. San Diego, CA, June 2023. DOI: 10.2514/6.2023-4018.
- [28] Vladimir Balepin, Christopher Osello, and Christopher Snyder. "NO_x Emission Reduction in Commercial Jets through Water Injection". In: *38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. Indianapolis, IN, July 2002. DOI: 10.2514/6.2002-3623.

- [29] Christos Mourouzidis et al. "Water injection on aircraft engines: a performance, emissions and economic study". In: *ISABE 22nd International Symposium on Air Breathing Engines*. Phoenix, AZ, Oct. 2015.
- [30] D. Alejandro Block Novelo et al. "Experimental investigation of gas turbine compressor water injection for NO_x emission reductions". In: *Energy* 176 (June 2019), pp. 235–248. DOI: 10.1016/j.energy.2019.03.187.
- [31] D. W. Bahr and T. F. Lyon. "NO_x Abatement via Water Injection in Aircraft-Derivative Turbine Engines". In: *Proceedings of the ASME 1984 International Gas Turbine Conference and Exhibit*. Vol. 3: Coal, Biomass and Alternative Fuels; Combustion and Fuels; Oil and Gas Applications; Cycle Innovations. Turbo Expo: Power for Land, Sea, and Air. Amsterdam, The Netherlands, June 1984, V003T06A014. DOI: 10.1115/84-GT-103.
- [32] Philip P. Walsh and Paul Fletcher. *An Introduction to Combustion: Concepts and Applications*. 2nd ed. Blackwell Science, 2004. ISBN: 9780632064342. DOI: 10.1002/9780470774533.
- [33] Joaquim R. R. A. Martins and Andrew Ning. *Engineering Design Optimization*. Cambridge, UK: Cambridge University Press, Jan. 2022. ISBN: 9781108833417. DOI: 10.1017/9781108980647. URL: <https://mdobook.github.io>.
- [34] Justin S. Gray et al. "OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization". In: *Structural and Multidisciplinary Optimization* 59.4 (Apr. 2019), pp. 1075–1104. DOI: 10.1007/s00158-019-02211-z.
- [35] John T. Hwang and Joaquim R. R. A. Martins. "A computational architecture for coupling heterogeneous numerical models and computing coupled derivatives". In: *ACM Transactions on Mathematical Software* 44.4 (June 2018), Article 37. DOI: 10.1145/3182393.
- [36] Joaquim R. R. A. Martins, Peter Sturdza, and Juan J. Alonso. "The Complex-Step Derivative Approximation". In: *ACM Transactions on Mathematical Software* 29.3 (Sept. 2003), pp. 245–262. DOI: 10.1145/838250.838251.
- [37] John K. Lytle. *The Numerical Propulsion System Simulation: An Overview*. Tech. rep. NASA/TM-2000-209915. NASA Glenn Research Center, June 2000.
- [38] David G. Goodwin et al. *Cantera: An Object-oriented Software Toolkit for Chemical Kinetics, Thermodynamics, and Transport Processes*. Version 3.0.0. 2023. DOI: 10.5281/zenodo.8137090. URL: <https://www.cantera.org>.
- [39] Arvind Gangoli Rao, Ivan Langella, and Dirk Roekaerts. *Hydrogen in Aero Engines: Opportunities and Challenges*. AE4262 Combustion for Propulsion and Power Technologies Lecture. 2022.
- [40] Gregorio Corchero and José Luis Montañés. "An approach to the use of hydrogen for commercial aircraft engines". In: *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering* 219.1 (2005), pp. 35–44. DOI: 10.1243/095441005X9139.
- [41] Vishal Sethi. *Hydrogen and Decarbonisation Workshop: Hydrogen and Low NO_x*. ISABE Conference. Sept. 2022. URL: <https://conference.isabe.org/content/scnotes/bobby/gtc13.pdf> (visited on 06/28/2024).
- [42] Y. B. Zeldovich, D. Frank-Kamenetskii, and P. Sadovnikov. "Formation of nitric oxide in premixed hydrocarbon flames". In: *Academy of Sciences of the USSR* (1947).
- [43] C. P. Fenimore. "Formation of nitric oxide in premixed hydrocarbon flames". In: *Symposium (International) on Combustion* 13.1 (1971), pp. 373–380. ISSN: 0082-0784. DOI: 10.1016/S0082-0784(71)80040-1.
- [44] A. N. Hayhurst and I. M. Vince. "Nitric oxide formation from N₂ in flames: The importance of 'prompt' NO". In: *Progress in Energy and Combustion Science* 6.1 (1980), pp. 35–51. ISSN: 0360-1285. DOI: 10.1016/0360-1285(80)90014-3.
- [45] James A. Miller, Michael J. Pilling, and Jürgen Troe. "Unravelling combustion mechanisms through a quantitative understanding of elementary reactions". In: *Proceedings of the Combustion Institute* 30.1 (2005), pp. 43–88. ISSN: 1540-7489. DOI: 10.1016/j.proci.2004.08.281.
- [46] Lyudmila V. Moskaleva and Ming-Chang Lin. "The spin-conserved reaction CH+N₂→H+NCN: A major pathway to prompt NO studied by quantum/statistical theory calculations and kinetic modeling of rate constant". In: *Proceedings of the Combustion Institute* 28.2 (2000), pp. 2393–2401. ISSN: 1540-7489. DOI: 10.1016/S0082-0784(00)80652-9.
- [47] P. C. Malte and D. T. Pratt. "Measurement of atomic oxygen and nitrogen oxides in jet-stirred combustion". In: *Symposium (International) on Combustion* 15.1 (Aug. 1975), pp. 1061–1070. DOI: 10.1016/S0082-0784(75)80371-7.
- [48] Tim C. Lieuwen and Vigor Yang, eds. *Gas Turbine Emissions*. 1st ed. Cambridge University Press, 2013. ISBN: 9781139015462. DOI: 10.1017/CB09781139015462.
- [49] James A. Miller, M. C. Branch, and Robert J. Kee. "A chemical kinetic model for the selective reduction of nitric oxide by ammonia". In: *Combustion and Flame* 43 (Dec. 1981), pp. 81–98. DOI: 10.1016/0010-2180(81)90008-0.
- [50] Richard K. Lyon. "Method for the reduction of the concentration of NO in combustion effluents using ammonia". US3900554A. Aug. 1975. URL: <https://patents.google.com/patent/US3900554A/>.
- [51] Richard K. Lyon. "The NH₃-NO-O₂ reaction". In: *International Journal of Chemical Kinetics* 8.2 (Mar. 1976), pp. 315–318. DOI: 10.1002/kin.550080213.
- [52] James A. Miller. "Theory and modeling in combustion chemistry". In: *Symposium (International) on Combustion* 26.1 (July 1996), pp. 461–480. DOI: 10.1016/S0082-0784(96)80249-9.
- [53] Joseph W. Bozzelli and Anthony M. Dean. "O + NNH: A possible new route for NO_x formation in flames". In: *International Journal of Chemical Kinetics* 27.11 (Nov. 1995), pp. 1097–1109. DOI: 10.1002/kin.550271107.
- [54] James A. Miller and Craig T. Bowman. "Mechanism and modeling of nitrogen chemistry in combustion". In: *Progress in Energy and Combustion Science* 15.4 (1989), pp. 287–338. ISSN: 0360-1285. DOI: 10.1016/0360-1285(89)90017-8.
- [55] Mauro A. Galbiati et al. "Mild combustion for fuel-NO_x reduction". In: *Combustion Science and Technology* 176.7 (July 2004), pp. 1035–1054. DOI: 10.1080/00102200490426424.
- [56] J. L. Toof. "A Model for the Prediction of Thermal, Prompt, and Fuel NO_x Emissions From Combustion Turbines". In: *Journal of Engineering for Gas Turbines and Power* 108.2 (Apr. 1986), pp. 340–347. ISSN: 0742-4795. DOI: 10.1115/1.3239909.
- [57] Jost O. L. Wendt. "Mechanisms Governing the Formation and Destruction of NO_x and Other Nitrogenous Species in Low NO_x Coal Combustion Systems". In: *Combustion Science and Technology* 108.4-6 (June 1995), pp. 323–344. DOI: 10.1080/00102209508960405.
- [58] L. D. Smoot, S. C. Hill, and H. Xu. "NO_x control through reburning". In: *Progress in Energy and Combustion Science* 24.5 (Oct. 1998), pp. 385–408. ISSN: 0360-1285. DOI: 10.1016/S0360-1285(97)00022-1.

- [59] André Nicolle and Philippe Dagaut. "Occurrence of NO-reburning in MILD combustion evidenced via chemical kinetic modeling". In: *Fuel* 85.17 (Dec. 2006), pp. 2469–2478. ISSN: 0016-2361. DOI: 10.1016/j.fuel.2006.05.021.
- [60] André A. V. Perpignan et al. "Emission Modeling of an Interturbine Burner Based on Flameless Combustion". In: *Energy Fuels* 32.1 (Jan. 2018), pp. 822–838. DOI: 10.1021/acs.energyfuels.7b02473.
- [61] D. Zhao et al. "Behavior and effect on NO_x formation of OH radical in methane-air diffusion flame with steam addition". In: *Combustion and Flame* 130.4 (Sept. 2002), pp. 352–360. ISSN: 0010-2180. DOI: 10.1016/S0010-2180(02)00385-1.
- [62] P. D. Norman et al. *Development of the technical basis for a new emissions parameter covering the whole aircraft operation: NEPAIR— Final Technical Report*. Tech. rep. NEPAIR/WP4/WPR/01. NEPAIR - European Commission, Sept. 2003.
- [63] N. Chandrasekaran and Abhijit Guha. "Study of Prediction Methods for NO_x Emission from Turbofan Engines". In: *Journal of Propulsion and Power* 28.1 (Jan. 2012), pp. 170–180. DOI: 10.2514/1.B34245.
- [64] Nicole Viola et al. "Analytical Formulations for Nitrogen Oxides Emissions Estimation of an Air Turbo-Rocket Engine Using Hydrogen". In: *Aerospace* 10.11 (Oct. 2023), p. 909. DOI: 10.3390/aerospace10110909.
- [65] Jocelyn Luche. "Obtention de modèles cinétiques réduits de combustion - Application à un mécanisme du kérosène". PhD thesis. Orléans, France: Université d'Orléans, Nov. 2003. URL: <https://theses.hal.science/tel-00636023/>.
- [66] Jocelyn Luche et al. "Reduction of large detailed mechanisms: Application to kerosene/air combustion". In: *Combustion Science and Technology* 176.11 (Nov. 2004), pp. 1935–1963. DOI: 10.1080/00102200490504571.
- [67] Philippe Dagaut et al. "Kerosene combustion at pressures up to 40 atm: Experimental study and detailed chemical kinetic modeling". In: *Symposium (International) on Combustion* 25.1 (July 1994), pp. 919–926. DOI: 10.1016/S0082-0784(06)80727-7.
- [68] Philippe Dagaut et al. "High pressure oxidation of normal decane and kerosene in dilute conditions from low to high temperature". In: *Journal de Chimie Physique* 92 (1995), pp. 47–76. DOI: 10.1051/jcp/1995920047.
- [69] Gregory P. Smith et al. *GRI-Mech 3.0*. URL: http://www.me.berkeley.edu/gri_mech/.
- [70] Qinghui Meng et al. "On the role of HNNO in NO_x formation". In: *Proceedings of the Combustion Institute* 39.1 (Jan. 2023), pp. 551–560. DOI: 10.1016/j.proci.2022.08.044.
- [71] Peter Glarborg et al. "Modeling nitrogen chemistry in combustion". In: *Progress in Energy and Combustion Science* 67 (Feb. 2018), pp. 31–68. DOI: 10.1016/j.peccs.2018.01.002.
- [72] Tommaso Capurso et al. "NO_x pathways in lean partially premixed swirling H₂-air turbulent flame". In: *Combustion and Flame* 248 (2023), p. 112581. ISSN: 0010-2180. DOI: 10.1016/j.combustflame.2022.112581.
- [73] International Civil Aviation Organization. *ICAO Aircraft Engine Emissions Databank*. June 2023. URL: <https://www.easa.europa.eu/en/domains/environment/icao-aircraft-engine-emissions-databank> (visited on 05/24/2024).
- [74] Pauli Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.
- [75] P. Saves et al. *SMT 2.0 Documentation v2.9.2: Kriging (KRG)*. Mar. 2025. URL: https://smt.readthedocs.io/en/v2.9.2/_src_docs/surrogate_models/gpr/krig.html (visited on 06/19/2025).
- [76] Arshad Afzal, Kwang-Yong Kim, and Jae-Won Seo. "Effects of Latin hypercube sampling on surrogate modeling and optimization". In: *International Journal of Fluid Machinery and Systems* 10.3 (Sept. 2017). DOI: 10.5293/IJFMS.2017.10.3.240.
- [77] Joaquim R. R. A. Martins and Andrew B. Lambe. "Multidisciplinary Design Optimization: A Survey of Architectures". In: *AIAA Journal* 51.9 (Sept. 2013), pp. 2049–2075. DOI: 10.2514/1.J051895.
- [78] Neil Wu et al. "pyOptSparse: A Python framework for large-scale constrained nonlinear optimization of sparse systems". In: *Journal of Open Source Software* 5.54 (Oct. 2020), p. 2564. DOI: 10.21105/joss.02564.



Sensitivity Analysis of a Cantera Reactor

This appendix covers several relevant topics related to the behaviour of the Cantera models, which help understand some of the approaches taken in this work. The first section covers the behaviour of the ignition pulse used to ignite the reactors, and the second section provides a sensitivity analysis of a Cantera reactor for the different input variables used in the models.

A.1. Ignition of the Mixture and Effects of the Igniter Pulse Shape

When analysing a combustor in Cantera, combustion ignition will not happen spontaneously unless the temperature of the reactants is above the autoignition temperature (which is not the case for many of the analysed conditions). Thus, combustion needs to be artificially triggered in the simulation.

The preferred method for this is to inject a pulse of hydrogen atoms into the combustors. These atoms, which are highly reacting free radicals, trigger the combustion reactions in the mechanism and lead to an increase in temperature. If the amount of hydrogen atoms injected is enough, the temperature will rise enough such that combustion will continue happening even if the hydrogen injection is cut.

This hydrogen pulse is not present in the real combustor, it is an artificial method used in Cantera. Thus, if the injection of hydrogen atoms were continuous, it would affect the steady state of the simulation and thus the simulation results. However, as said before, the injection of hydrogen can be cut once the mixture has been ignited. After this, the simulation is converged to steady state, such that the effect of hydrogen injection disappears. For this, the mass flow of the igniter is set to be a time-dependent pulse.

The pulse is defined by two main considerations: the total amount of hydrogen injected and the width of the pulse over which it is injected. A higher amount of hydrogen will lead to better ignition results, and a narrower pulse will also lead to a higher successful ignition rate (for a narrower pulse, more hydrogen is present inside the reactor at any one time, so more of the reactor content reacts). Thus, the profile of the pulse needs to be chosen such that ignition is successfully achieved.

It is also important to verify that this artificial pulse does not affect the steady state results; the simulation needs to converge to steady state after the pulse is turned off. To prove that the effects of the pulse do




Figure A.1: The shape of the pulse affects the transient behaviour of the simulation, but does not affect the steady-state results as long as ignition is successfully achieved.

no affect the steady state results, a reactor with a fixed set of initial conditions (composition, temperature, pressure, etc.) was ignited with pulses with different hydrogen amounts and pulse widths. The results of this are shown in Figure A.1.

The results clearly show that the shape of the pulse affects the transient behaviour of the simulation significantly. However, it also shows that all the successfully ignited simulations converge to the same final value, so the shape of the ignition pulse does not have an effect on the converged result beyond causing ignition. The analysis also shows that if the pulse is not strong enough (too little total hydrogen mass, or too wide of a pulse), the combustion will not be sustained and the model will return to initial conditions.

A.2. Sensitivity of the Input Parameters

It is also important to understand the behaviour of a reactor with different input conditions. This information can help to understand the effects that making modification to the network will have, and it can also help to spot errors in the implementation of the network. A sensitivity analysis using a single perfectly stirred reactor was performed. The reaction mechanism used was Luche's mechanism [65], using n-decane ($nC_{10}H_{22}$) as fuel, and air as oxidiser. The basic conditions of the reactor are shown in Table A.1. Each of the input parameters was then varied independently to understand the effect on the reactor properties. In all cases the characteristics of the ignition pulse are kept constant, to understand the effect that the parameters have on successfully establishing ignition in the reactor.

Table A.1: Baseline conditions of the reactor used for the sensitivity analysis.

Parameter	Unit	Value
T	K	600
p	atm	1
\dot{m}_{air}	kg/s	1
φ	—	1
V	m ³	1

The first parameter analysed is the inlet temperature. As expected, the results in Figure A.2 show that an increase in inlet temperature results in an increase in exit temperature. The results also show that Cantera can capture self-ignition if the temperature is high enough. In those cases, the mixture ignites instantly without the need of an ignition pulse. However, the pulse was still present in those simulations for consistency, and the transient effect of the ignition pulse (a temperature rise centred around the 2-second mark, where the ignition pulse is centred) can be seen.




Figure A.2: Increasing inlet temperatures result in higher exit temperatures, and might trigger autoignition in the reactor.

The second parameter analysed is the inlet pressure. In this case, the igniter mass was also varied for each pressure level. This was done to should an example of the coupled effect of the igniter parameters and the input parameters on the transient behaviour of ignition. The results show that a higher inlet pressure results in a higher exit temperature, but with a more moderate effect compared to that seen with the inlet pressure.

An increase in pressure also hinders ignition of the mixture, again due to the higher mass in the reactor for a higher pressure.




Figure A.3: Higher pressures result in marginal gains in combustion temperature, but also hinder ignition due to the higher reactor mass.

The next parameter analysed is the mass flow. The results are shown in Figure A.4. Increasing the mass flow has a minimal but negative effect in the outlet temperature, with the temperature dropping by a few Kelvin. If the mass flow is increased further, it can make ignition unsuccessful.




Figure A.4: Increasing the mass flow has a very small negative effect on exit temperature, and can also negatively affect ignition.

The effect of changing the equivalence ratio is shown in Figure A.5. The behaviour of the results is as expected: the highest temperature is seen for equivalence ratios slightly above 1 (the stoichiometric ratio), decreasing for both higher and lower equivalence ratios.

The same results, including only the final temperature, are shown in Figure A.6, but presented in the traditional manner: temperature versus equivalence ratio. This graph shows the same behaviour more clearly, with the maximum temperature at an equivalence ratio slightly above stoichiometry.

The last parameter analysed is the combustor volume, with the results shown in Figure A.7. A very low combustor volume will result in a low residence time and consequently in incomplete combustion, leading to reduced exit temperatures. Increasing the volume causes the residence time to increase, thus allowing for more complete combustion, and resulting in the exit temperature approaching the adiabatic flame temperature. However, when the volume is increased beyond a certain point, ignition fails. This is due to the fact that the starting mass of the reactor is proportional to the volume (since T and p are fixed), so the amount of hydrogen injected for ignition becomes insufficient.




Figure A.5: The equivalence ratio behaves as expected, with the highest temperature being around stoichiometric conditions.




Figure A.6: The maximum temperature happens at an equivalence ratio slightly above 1, decreasing for both lower and higher equivalence ratios.




Figure A.7: Increasing the combustor volume results in higher residence times and therefore more complete combustion, but increasing the reactor volume (and thus its mass content) beyond a certain point can make ignition unsuccessful.

Hydrogen Engine with Water Recirculation

This appendix contains all the information and results regarding the hydrogen combustor with water recirculation. A hydrogen engine with water recirculation was not included in the final set of results, but was investigated during this work, so it is added here for completeness and for reference.

B.1. Extension of the Cantera Model for Water Recirculation

The model for a hydrogen combustor with water injection is based in the model from Section 4.3, to which water injection was added by adding a mass flow controller that connects a reservoir with 100% water vapour, and feeds it to the mixer reactor. The mass flow rate associated with the mass flow controller is left as an input to the model. The diagram with the architecture of the reactor is shown in Figure B.1.

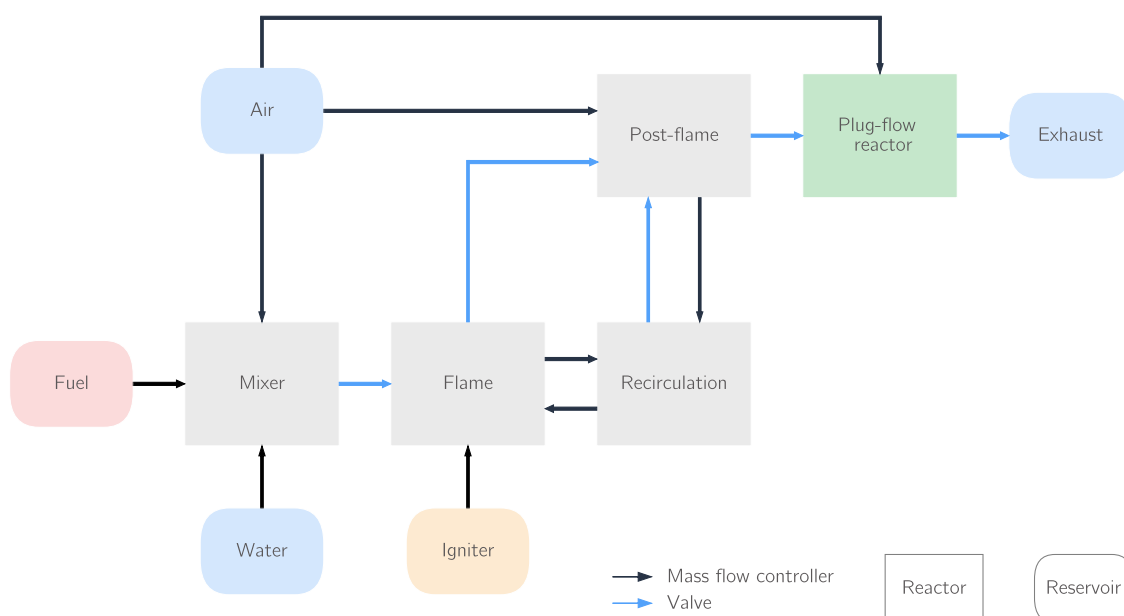


Figure B.1: The architecture of the CRN of the hydrogen combustor with water injection is the same as for the dry hydrogen combustor, with the addition of a water reservoir that feeds into the mixer reactor.


B.2. Surrogate Model

The background theory and the procedure for the generation of the surrogate models is described in Chapter 5, as well as the specifics regarding the sampling of the Cantera models. Table B.1 shows the bounds of the sampling variables for the hydrogen combustor model with water injection.

For this model, the error flattens at 18 points per variable (108 total points), thus this number of sampling points was selected. As was done for the other models, once the number of sampling points per variable was selected, a new set of points using Latin hypercube sampling was created for the model.

Table B.1: Bounds of the sampling intervals for each of the input functions of the H₂ surrogate model.

Parameter	Unit	Lower bound	Upper bound
T ₃	K	600	1000
p ₃	Pa	14e5	25e5
\dot{m}_{air}	kg/s	8	20
FAR	—	0.005	0.012
A	m ²	0.02	0.08
\dot{m}_{water}	kg/s	0	1.5

**Figure B.2:** The error of the H₂ model with water injection starts to flatten at 18 sampling points per variable (108 total sampling points).

B.3. Optimisation Results

The optimisation performed here is analogous to case *JW* described in Chapter 6, where the water extraction fraction is the same across all four design points. The results are shown in Table B.2.

Table B.2: Results of the optimisation of a hydrogen turbofan with water recirculation. A lower asterisk * represents an active lower bound or \geq constraint, and a high asterisk represents an active upper bound or \leq constraint. The values in square brackets are reported for completeness, but were not implemented as constraints or design variables.

		H₂ - Wet	
<i>Objective</i>	$\dot{m}_{fuel, CRZ}$	kg/s	0.10511
<i>Constraints</i>	EINO _{xCRZ}	g/kg	[0.6255]
	F _{TOC}	kN	25.8*
	D _{fan}	m	2.54*
<i>Design variables</i>	$\Pi_{fan, TOC}$	—	1.284
	$\Pi_{LPC, TOC}$	—	4*
	OPR _{TOC}	—	53.00
	T _{4, RTO}	K	1766.38
	T _{4, ratio}	—	0.9066
	V _{ratio, CRZ}	—	1.35*
	X _{H₂O}	—	0.0833

The combustor conditions that are inputs to the surrogate model are shown in Table B.3. The water extraction fraction is lower than for the Jet A engine. However, given that hydrogen engines have more water content in the exhaust gas, the absolute value of the extracted water mass flow is approximately twice

higher than that of the Jet A engine during cruise. The NO_x emissions are very low compared to the Jet A engines, similarly to what was observed in the dry engines. The combustor inlet temperature for this case is the highest among all analysed cases in this work, but the FAR ratio is also the lowest, which helps offset the rise in NO_x that a higher T₃ would entail. The NO_x emissions are lower than for the dry H₂ engine, which was not the case for the Jet A engine.

Table B.3: Inputs to the emissions surrogate models for each of the cases with water recirculation.

		H₂ - Wet
T ₃	K	927.88
p ₃	Pa	17.66
\dot{m}_{air}	kg/s	14.90
FAR	—	0.00705
ϕ	—	0.2397
A	m ²	0.0537
\dot{m}_{water}	kg/s	0.0736
EINO _x	g/kg	0.6255