A Novel Turbomachine for Decarbonisation of High-Temperature Chemical Reaction Processes and Fast Aerochemical Coupling



Dylan Austin Rubini

Supervisor: Professor Budimir Rosic

Department of Engineering Science

University of Oxford

This dissertation is submitted for the degree of

Doctor of Philosophy

St Hugh's College

November 2024

This thesis is dedicated to Alicja,
who carried me on this beautiful journey

Acknowledgements

Above all else, I am forever immensely grateful to my supervisor, Professor Budimir Rosic. His endless support, encouragement during dark times, guidance and wisdom throughout my DPhil have been unparalleled. Most importantly, he provided the perfect balance between freedom to explore, learn new things and make mistakes, while also offering insightful, creative and exceptionally innovative ideas to guide me precisely when needed them most. I am fortunate for Budimir's unwavering optimism that counterbalanced my recurring pessimism quite well.

I would also like to thank Professor Liping Xu for frequently sharing his unrivalled knowledge about turbomachinery and engineering, always forcing me to learn more just to keep up. I thank Craig for always responding quickly to technical issues with my computer.

It has been an honour working with the fantastic friends in Budimir's lab group—Oguzhan, Nikolas, Simla, Jamie, Mark and many others. Thank you Nikolas for always spurring me on and Oguzhan for teaching me about interesting things during numerous cigarette breaks.

I am grateful to Coolbrook, particularly Dr Elina Nauha, Dr Tuomous Ouni and Harri Johannesdahl, for their support and fruitful discussions.

Finally, thank you to all my friends and family for always keeping life interesting. And Alicja, you are the light in my life, I couldn't have done this without you.

Abstract

Decarbonising *high-temperature* industrial processes—steel, cement, and petrochemicals—poses one of the greatest challenges in achieving net-zero emissions. To address this, this thesis introduces a revolutionary turbomachinery concept, the turbo-reactor, which replaces conventional surface heat exchange within furnaces by directly transferring mechanical energy to the fluid through a renewably powered, electric-motor-driven system. This fundamental shift in the energy transfer mechanism offers substantial improvements in power density, scalability, energy efficiency, operability, and reaction performance compared to both traditional technologies and alternative decarbonisation strategies.

Using a combination of high- and low-fidelity computations, this work presents the working principles of the elemental stage design and investigates the uniquely complex aerothermodynamics within both axial and regenerative turbo-reactor architectures. The feasibility of an ultra-high-work-coefficient stage design complemented by rapid dissipation is successfully demonstrated. It is concluded that this universal stage design philosophy can be used effectively across a range of feedstocks and chemical reaction states.

For many applications of the turbo-reactor in the chemical/petrochemical industries, energy is directly supplied to the process gas to drive an endothermic chemical reaction. Therefore, this opens up a new design space for optimising the reaction performance through careful aerodynamic design. The primary objective is to design a reaction-efficient temperature profile by balancing the reaction heat absorption with fine-tuned control of the rate of work input and energy dissipation. Consequently, the physical flow mechanisms responsible for the energy dissipation process are explored in detail. Then, it is outlined how the temperature profile can be tailored by adjusting the distance between stages. From an aerodynamic performance perspective, the minimum distance between stages is determined from high-fidelity computations. This sets a lower bound on the range of control available to the aerochemical designer. Using simple 0D chemical reaction simulations, it is established that there is scope for optimising the temperature profile to improve reaction performance. However, this is difficult to achieve in practise since 3D viscous reacting flow simulations with detailed and accurate kinetic models are prohibitively costly.

Traditional state-of-the-art acceleration strategies for reactive flows are insufficient to bring aerochemical simulations into the design optimisation loop. To address this, a new multi-fidelity machine-learning-enhanced efficiently coupled aerochemical modelling methodology, called `ChemZIP`, is developed. It is demonstrated that `ChemZIP` is nearly two orders of magnitude faster than state-of-the-art acceleration approaches while maintaining accuracy within 10%. For the first time, this enables aerochemical numerical simulations to be brought into the aerodynamic design optimisation outer loop. This will lead to more efficient designs that can be developed quickly and with a reduced reliance on costly physical experiments.

Publications

1. Rubini, D., Xu, L., Rosic, B., and Johannesdahl, H. "A New Turbomachine for Clean and Sustainable Hydrocarbon Cracking". In: *ASME J. Eng. Gas Turbines Power* 144.2 (Dec. 2021), pp. 021024.*

2. Rubini, D., Karefyllidis, N., Xu, L., Rosic, B. and Johannesdahl, H. "A New Robust Regenerative Turbo-Reactor Concept for Clean Hydrocarbon Cracking". In: *J. Glob. Power Propuls. Soc.* 6 (2022), pp. 135–150.†

3. Rubini, D., Karefyllidis, N., Xu, L., Rosic, B., and Johannesdahl, H. "Accelerating the Development of a New Turbomachinery Concept in an Environment With Limited Resources and Experimental Data: Challenges". In: Volume 10D: *Turbomachinery - Multidisciplinary Design Approaches, Optimization, and Uncertainty Quantification; Turbomachinery General Interest; Unsteady Flows in Turbomachinery* (Rotterdam, Netherlands, 2022). Turbo Expo: Power for Land, Sea, and Air. June 2022, V10DT36A001.†

4. Karefyllidis, N., Rubini, D., Rosic, B., Xu, L., and Purola, V. "A Novel Axial Energy-Imparting Turbomachine for High-Enthalpy Gas Heating: Robustness of the Aerodynamic Design". In: *ASME J. Turbomach.* 146.3 (Nov. 2023), pp. 031005.†

5. Rubini, D., Karefyllidis, N., Rosic, B., Xu, L., and Nauha, E. "Decarbonisation of High-Temperature Endothermic Chemical Reaction Processes Using a Novel Turbomachine: Robustness of the Concept to Feed Variability". In: *J. Glob. Power Propuls. Soc.* 8 (2024), pp. 111–126.†

6. Rubini, D., Rosic, B., and Xu, L. "Energy Conversion Mechanisms in Turbomachines for High-Temperature Endothermic Reactions: Redefining Energy Quality". In: *Appl. Therm. Eng.* 258 (2025), pp. 124566. †

*__Not included__ in this thesis
†__Included__ in this thesis

Table of Contents

List of Figures **v**

List of Tables **x**

Nomenclature **xi**

I Aerothermodynamics **1**

1 Introduction **2**

1.1 Decarbonising High-Temperature Industrial Heat 2

1.1.1 Direct Electrification of High-Temperature Industrial Heat with a New Turbomachine 4

1.2 Research Questions and Thesis Structure 8

2 A New Turbomachinery Concept for High-Temperature Processes **12**

2.1 A Review of Concepts to Decarbonise Chemical Processes 12

2.1.1 Conventional Hydrocarbon Cracking Process 13

2.1.2 Power-to-X: Carbon-Neutral Fuels 16

2.1.3 Electric Furnaces 20

2.1.4 Electrochemical 24

2.1.5 Outlook 24

2.2 A New Turbomachine for Clean Pyrolysis 27

2.2.1 Introduction 27

2.2.2 Perspective 49

3 Redefining Energy Quality for the Turbo-Reactor **52**

3.1 Energy Conversion Mechanisms 52

3.1.1 Perspective 69

4 Robustness of the Turbo-Reactor Aerodynamics 70
4.1 Flexibility of the Axial Multi-Stage Arrangement . . . 70
4.1.1 Chemically Tuning the Multi-Stage Arrangement . . . 83
4.1.2 Summary . . . 90
4.2 Robustness of the Concept to Feed Variability . . . 91
4.3 Perspective . . . 110

5 The Design Ecosystem: New Requirements 111
5.1 Accelerating the Turbo-Reactor Development . . . 111
5.1.1 Summary . . . 130

II Aerothermochemistry 131

6 Modelling Aerochemical Interactions 132
6.1 Introduction . . . 132
6.1.1 Importance of 3-D Viscous Aerochemical Interactions . . . 134
6.1.2 Challenges of Reacting Flow Modelling . . . 138
6.2 Fully-Coupled Detailed Reacting Flow Model . . . 139
6.2.1 Conservation Equations . . . 139
6.2.2 Mass-Action Kinetics and the Arrhenius Rate Law . . . 141
6.2.3 Physical Modelling Assumptions . . . 142
6.3 Accelerated Aerochemical Modelling Literature Survey . . . 146
6.3.1 *Physics-Driven* Strategies . . . 146
6.3.2 *Data-Driven* Strategies . . . 152
6.4 Summary . . . 156

7 Efficiently Coupled Aerochemical Modelling Methodology 158
7.1 Why Machine Learning? . . . 158
7.2 Training Data Generation . . . 160
7.2.1 Sample Space Generation . . . 162
7.2.2 Constant-Pressure 1-D PFR Simulations . . . 163
7.2.3 Thermochemical Phase Space . . . 165
7.2.4 The Training Database . . . 167
7.3 Neural Network Architecture . . . 167
7.3.1 Linear or Non-Linear Compression? . . . 168
7.3.2 Non-Linear Reconstruction . . . 171
7.3.3 Latent Space: Thermochemical Dynamic Response . . . 171
7.3.4 The Loss Function . . . 173
7.3.5 Data Pre-Processing Pipeline . . . 176
7.3.6 Error Metrics . . . 177

7.4 Latent Phase Space Manifold . . . 177
7.5 Coupling `ChemZIP` with a Flow Solver . . . 178
7.5.1 TBLOCK–`ChemZIP` Interface . . . 180
7.6 Summary . . . 181

8 Verification, Benchmarking and Proof of Concept 183
8.1 Introduction . . . 183
8.2 Verification: Error Statistics . . . 184
8.2.1 A Priori Input-Output Reconstruction . . . 185
8.2.2 A Posteriori Multi-Step 1-D PFR Predictions . . . 187
8.3 Verification: Three-Dimensional Heated Duct Flow . . . 193
8.3.1 Numerical Method and Boundary Conditions . . . 193
8.3.2 Computational Acceleration . . . 196
8.3.3 Predictive Accuracy . . . 200
8.4 Proof of Concept: Multi-Stage Turbo-Reactor . . . 205
8.5 Summary . . . 207

9 Conclusions and Future Work 209
9.1 Discussion of Major Contributions . . . 209
9.1.1 The New Turbo-Reactor Concept . . . 209
9.1.2 Energy Conversion Mechanisms . . . 210
9.1.3 Robustness and Flexibility of the Turbo-Reactor . . . 210
9.1.4 Accelerating Aerochemical Computations . . . 212
9.2 Recommendations for Future Work – Aerochemistry . . . 214
9.2.1 Network Architecture and Learning Algorithm . . . 214
9.2.2 Physical Modelling Assumptions . . . 215
9.2.3 Aerodynamic–Coke Interaction . . . 217
9.2.4 Catalytic Reaction Modelling . . . 217
9.2.5 Future Applications . . . 219
9.3 Final Remarks . . . 219

Appendices

A Coke Formation Model 221

B Detailed Neural Network Architecture 223
B.1 Detailed Network Structure . . . 223
B.2 Details of the Training Procedure . . . 224

References 225

List of Figures

1.1 Breakdown of global energy demand by sector [4] 2
1.2 Breakdown of global greenhouse gas emissions by sector [7] 4
1.3 Electricity price and carbon tax necessary for industrial electrification of heat to become cost competitive [9] 5
1.4 Schematic illustration of historical and 40-years projections for the levelised cost of energy (LOCE) including storage and carbon tax (based on data from [10, 11]) 6
1.5 Cost breakdown for high-value chemicals production [13] 7
1.6 A turbo-reactor render [14] and some example industries that it could decarbonise 8

2.1 Conventional steam cracking plant schematic (adapted from [24, 25]) . . . 14
2.2 Coke accumulation within a conventional ethylene furnace coil (Refs. [26, 27]) 15
2.3 Biomethane production process [40] 18
2.4 Green hydrogen production via water electrolysis [44] 19
2.5 Joule heated electrified steam methane reforming reaction for H_2 production [56] 21
2.6 Magnetic induction heating for thermochemical processes using metamaterial-based foams [62] 22
2.7 Schematic of a dielectric discharge non-thermal plasma chemical reactor [69] 23
2.8 Global map of renewable energy resource potential [76] 26
2.9 A schematic of the shockwave reactor used for the pyrolysis of hydrocarbons (reworked based on [84, 85]) 28

4.1 1-D plug-flow reactor approximation of the multi-stage turbo-reactor used for chemical kinetic calculations 84
4.2 For an increasing number of stages $N_{\text{stages}} \in [10, 16]$: $(\boldsymbol{a})$ the Pareto front of reaction performance $(\boldsymbol{b})$ the cumulative work input distribution and $(\boldsymbol{c})$ the temperature distribution 87
4.3 Effect of increasing the maximum allowable temperature level constraint on the Pareto front of reaction performance 88

4.4 For uniform and non-uniform enthalpy input cases: **(*a*)** reaction performance Pareto fronts **(*b*)** cumulative work input distribution and **(*c*)** temperature distribution 89

6.1 Chemical kinetic calculations of a 1D PFR demonstrating the risks of using a simplified (small) kinetic model compared to a detailed, comprehensive and accurate (large) model 133
6.2 Streamwise distribution of characteristic chemical time-scales plotted on a log-scale for an propane feed calculated from 1-D PFR chemical kinetic calculations. The characteristic chemical time scales are calculated using method outlined in Ref. [101]. Time-scales longer than 1000 s (*i.e.,* non-reacting species) are eliminated 136
6.3 Steady, inviscid, adiabatic 1-D compressible flow calculations across a shockwave for a steam-propane gas mixture that is (1) non-reacting (2) in chemical equilibrium and (3) in chemical non-equilibrium. The conditions upstream of the shock are representative of those of the turbo-reactor. The governing equations can be found in Anderson [102] 137
6.4 Reaction path diagram for a propane pyrolysis chemical mechanism with 65 species and 1700 reactions 138
6.5 Comparison of ignition times obtained with a detailed iso-octane mechanism (850 species) and skeletal mechanisms of various sizes (reproduced from [121]) 148
6.6 Schematic illustration of a low-dimensional manifold in composition space 150
6.7 Verification of FGM against detailed chemistry for a 2D premixed laminar methane flame (reproduced from [134]) 152
6.8 Artificial neural network schematic illustration 154
6.9 Improving low-dimensional projected manifold quality (reproduced from [158]) 156

7.1 Schematic illustration of the reaction response to representative temperature gradients in the turbo-reactor compared to those in a tubular pipe of a conventional furnace 161
7.2 Schematic of the training data generation workflow for the `ChemZIP` methodology 162
7.3 Typical temperature gradients $\frac{DT}{Dt} = \vec{\mathbf{u}} \cdot \nabla T$ ($\mathrm{K\,\mu s^{-1}}$) seen by a fluid parcel travelling along the turbo-reactor gas path computed using non-reacting CFD 163
7.4 A selection of 10 (out of 10,000) example training data 1-D PFR reactor simulations showing **(*a*)** the heat flow **(*b*)** the reaction heat absorption rate on a log scale **(*c*)** the dry propane yield and **(*d*)** the propane net production rate on a log scale 164

7.5 Schematic illustration of time step clustering and sub-sampling (plotted on a linear scale). Note, to help visualisation, only a fraction of the time steps are plotted 165
7.6 Scatter plot showing a projection of the high-dimensional thermochemical phase space for the training data, including noise augmented samples. Note that the number of training data points is downsampled by a factor of 20 and a log-scale used to help visualisation 166
7.7 Schematic illustration showing the `ChemZIP` chemical autoencoder NN architecture for compressing the thermochemistry and coupling with a 3-D viscous CFD flow solver 168
7.8 Schematic of isolated encoder-decoder architecture extracted from the full `ChemZIP` workflow in Fig. 7.7 170
7.9 **(*a*)** Schematic illustration of the power transform pre-processing function and **(*b*)** schematic illustration of various loss functions 174
7.10 A histogram of the training database for the reaction heat absorption rate plotted on **(*a*)** a linear scale and **(*b*)** a logarithmic scale 175
7.11 Latent space manifold for the training database constructed from Z_1, Z_2 &T and coloured by the reaction heat absorption rate $\dot{\omega}_T$ as the dependent variable (on a log-scale) 178
7.12 Flow chart of the coupling process between the CFD flow solver and the trained `ChemZIP` model written using the TensorFlow API 180
7.13 A verification of pressure distribution prediction on a highly-loaded gas turbine nozzle guide vane from U-Tblock against Fluent 181

8.1 A priori (*i.e.,* no time integration) parity plot (500,000 data points) comparing reconstructed dry yields predicted by `ChemZIP` against the ground truth FOM data for 5 species **(*a*)**–**(*e*)** 185
8.2 A priori (*i.e., without* time integration) parity plot (500,000 data points) comparing reconstructed `ChemZIP` predictions against the ground truth FOM data for **(*a*)**–**(*d*)** the net production rates **(*e*)** the heat absorption rate (on a log scale) and **(*f*)** the heat capacity 186
8.3 1D PFR time integrated solution history for **(*a*)** temperature **(*b*)**–**(*c*)** major/minor species dry yields **(*d*)** the reaction heat absorption rate and **(*e*)**–**(*f*)** the major/minor species net production rates. Note that the circular markers denoting the FOM are significantly downsampled for clarity 188
8.4 A posteriori (*i.e., with* time integration) parity plot (500,000 data points) comparing reconstructed `ChemZIP` predictions against the ground truth (FOM) data for **(*a*)**–**(*b*)** the propane/ethylene dry yields **(*c*)** the heat absorption rate (on a log scale) and **(*f*)**–**(*e*)** the propane/ethylene production rates 189

8.5 Histograms for the relative error $(x) = |(x - \widehat{x})/x|$ between `ChemZIP` and the ground truth FOM for time-integrated (a posteriori) and non-time-integrated (a priori) test data for **(*a*)** the reaction heat absorption rate and **(*b*)** all the species mass fractions. The abscissa is on a log scale . . . 190
8.6 2D joint and marginal distributions for the relative prediction errors and absolute value for the reaction heat absorption rate (500,000 test data points). Both axes are on a log scale . . . 191
8.7 Comparing `ChemZIP` and the FOM predictions under parametric design space exploration for the heat input, showing **(*a*)** the high-value primary product (ethylene + propylene + 1 3-butadiene + hydrogen + 1-butene + 2-butene) yield Y_{PP} and steady state coking rate distributions and **(*b*)** the coking rate and primary product yields at the outlet as a function of heat input . . . 192
8.8 Heated duct test case domain and computational mesh (totalling 1.1 million nodes) . . . 194
8.9 Convergence history for propane yield, which is the slowest variable to reach steady state ($\dot{Q} = 150\,\mathrm{kW\,m^{-2}}$) . . . 195
8.10 A comparison of static temperature and Mach number contours on an mid x-y plane cross section computed from a non-reacting flow simulation ($\dot{Q} = 100\,\mathrm{kW\,m^{-2}}$) for both solvers . . . 196
8.11 Computational acceleration achieved by `ChemZIP` compared with direct integration (DI) in FLUENT and a state-of-the-art acceleration (Accel.) strategy that combines DACM + DCZ . . . 197
8.12 Breakdown of the normalised CPU time per CFD iteration, comparing `ChemZIP` with FLUENT for a 65-species mechanism . . . 198
8.13 For a 65-species propane mechanism: **(*a*)** a comparison of the average numerical stiffness for the chemistry (on a log scale) and **(*b*)** the normalised single-core CPU time required to time integrate the `ChemZIP` model with either an stiff implicit integrator or an explicit integrator . . . 199
8.14 For a 65-species propane mechanism: **(*a*)** a comparison of the streamwise variation in minimum characteristic chemical time-scale between `ChemZIP` and the FOM (on a log scale) and **(*b*)** comparison of single-core CPU time to evaluate the chemical and energy source terms as a function of number of mesh cells in a computational batch (on a log-log scale) . . . 200
8.15 Verification of the reconstructed temperature and dry yields under two heat loads . . . 201
8.16 Verification of the reconstructed the reaction source terms under two heat loads . . . 202
8.17 A comparison of the propane yield distributions for `ChemZIP` both with and without species diffusion activated ($\dot{Q} = 100\,\mathrm{kW\,m^{-2}}$) . . . 203

8.18 Net species production rate distributions on an z-y plane cross section ($\dot{Q} = 100\,\mathrm{kW\,m^{-2}}$) 204
8.19 A projection of the latent phase space manifolds for all points in the training database and for all solution points from the heated duct ($\dot{Q} = 150\,\mathrm{kW\,m^{-2}}$) CFD simulation (shown in red) 205
8.20 Time-averaged **(*a*)** change in major and minor species yields relative to the inlet and **(*b*)** ethylene net production rates (on log scale) at midspan across a two-stage turbo-reactor 206
8.21 Time-averaged steady state coke formation rate distribution in units of µm per day on the suction and pressure surfaces on the diffuser blade in the turbo-reactor 207

9.1 Illustration of coke growth on the blade surfaces 218

List of Tables

2.1 A summary of conversion efficiency and cost of energy for different methods to decarbonise high-temperature industrial heat [46, 73, 74] 25

4.1 Optimisation parameter space and constraints 86

7.1 Sample space parameter ranges for each heat flux profile control point . . 163

8.1 Error metrics for the training dataset and the test dataset for both single-step (a priori) and multi-step (a posteriori) predictions 187
8.2 Boundary conditions for heated duct 194

B.1 Detailed neural network architecture for the encoder 223
B.2 Detailed neural network architecture for the decoder 223
B.3 Detailed neural network architecture for the dynamics mapping 223
B.4 Hyperparameters of the learning algorithm 224
B.5 Loss function weights from Eq. (7.7) 224

Nomenclature

Roman Symbols

A	Area	m
c_p	Isobaric specific heat capacity	$\mathrm{J\,kg^{-1}\,K^{-1}}$
C_x	Diffuser axial chord length	m
$C_{\Psi_{\mathrm{TS}}}$	Total-to-static pressure coefficient	–
$\bar{D}_k$	Molecular diffusivity of k^{th} species	$\mathrm{m^2\,s^{-1}}$
$\boldsymbol{D}_{\mathrm{eff}}$	Total diffusivities matrix $\in \mathbb{R}^{K\times K}$	$\mathrm{m^2\,s^{-1}}$
$\boldsymbol{D}_Z$	Meta species diffusivities matrix $\in \mathbb{R}^{M\times M}$	–
E	Total non-chemical energy	$\mathrm{J\,kg^{-1}}$
f	Frequency of a disturbance	Hz
$\mathscr{F}_{\mathrm{NN,Z}}$	Latent space dynamic response NN approximator	–
$\mathscr{F}_{\mathrm{NN}}$	Neural network function approximator	–
h_0	Stagnation (or total) enthalpy	$\mathrm{J\,kg^{-1}}$
h_k	Partial enthalpy of the k^{th} species	$\mathrm{J\,kg^{-1}}$
$\Delta h^\circ_{f,k}$	Formation enthalpies	$\mathrm{J\,mol^{-1}}$
$\mathbf{J}$	Species net production rates Jacobian matrix $\in \mathbb{R}^{K\times K}$	$\mathrm{s^{-1}}$
$\mathbf{J}_\phi$	Encoder Jacobian matrix $\in \mathbb{R}^{M\times K}$	–
$\mathbf{J}_\psi$	Decoder Jacobian matrix $\in \mathbb{R}^{K\times M}$	–
K	Number of species	–
M	Number of latent meta species	–
M_k	Molecular weight for the k^{th} species	$\mathrm{kg\,mol^{-1}}$
N	Number data points per variable in the training database	–
N_r	Number of reactions	–
p	Static pressure	Pa
$\dot{Q}$	Imposed heat flux	$\mathrm{kW\,m^{-2}}$
R	Specific gas constant	$\mathrm{J\,kg^{-1}\,K^{-1}}$
R^2	Coefficient of determination (goodness of fit metric)	–

R_u	Universal gas constant	$\mathrm{J\,mol^{-1}\,K^{-1}}$
T	Static temperature	K
t	Time	s
$\vec{\mathbf{u}}$	Velocity vector	$\mathrm{m\,s^{-1}}$
U	Mean blade speed	$\mathrm{m\,s^{-1}}$
V	Absolute velocity	$\mathrm{m\,s^{-1}}$
$\mathbf{W}$	Output dynamic response training data matrix $\in \mathbb{R}^{N\times K+4}$	–
$\mathbf{X}$	Input space training data matrix $\in \mathbb{R}^{N\times K+2}$	–
X_k	Mole fraction of the k^{th} species	%
$\boldsymbol{Y}$	Mass fraction vector $\in \mathbb{R}^{K}$	%
Y_k	Mass fraction of the k^{th} species	%
$\boldsymbol{Z}$	Latent space meta species mass fractions vector $\in \mathbb{R}^{M}$	–
Z_k	Mass fraction for the m^{th} latent meta species	–

Greek Symbols

α_T	Molecular thermal diffusivity	$\mathrm{m^2\,s^{-1}}$
δ_p	Thermal penetration depth	m
$\eta_{\mathrm{conversion}}$	Energy conversion efficiency	%
γ	Ratio of heat capacities	–
λ	Molecular/laminar conductivity	$\mathrm{W\,m^{-1}\,K^{-1}}$
μ	Molecular/laminar viscosity	$\mathrm{Pa\,s}$
μ_t	Turbulent viscosity	$\mathrm{Pa\,s}$
$\dot{\boldsymbol{\omega}}$	Species net production rate vector $\in \mathbb{R}^{K}$	$\mathrm{s^{-1}}$
$\dot{\omega}_k$	Net production rate of k^{th} species	$\mathrm{s^{-1}}$
$\dot{\omega}_T$	Reaction heat release/absorption rate	$\mathrm{W\,m^{-3}}$
$\dot{\boldsymbol{\omega}}_Z$	Net production rate vector for the latent meta species $\in \mathbb{R}^{K}$	–
$\dot{\omega}_{Z_m}$	Net production rate of the m^{th} latent meta species	–
ρ	Density	$\mathrm{kg\,m^{-3}}$
σ_T^2	Temperature scalar variance	$\mathrm{K^2}$
τ_{c_k}	Characteristic chemical time-scale	s
τ_f	Characteristic fluid time-scale	s

Non-Dimensional Groups

$Da_{\mathrm{I},k} = \frac{\tau_f}{\tau_{c_k}}$	Damköhler number for the k^{th} species
Le_k	Lewis number for k^{th} species

Ma	Mach number
Sc_t	Turbulent Schmidt number

Superscripts

$\bar{\square}$	Mixture averaging
$\overline{\square}$	(Time/Reynolds) averaging
$\square'$	Fluctuating component
$\widehat{\square}$	Reconstructed/predicted quantity

Subscripts

$\square_0$	Stagnation quantity
$\square_j$	Reaction index
$\square_k$	Species index
$\square_m$	Latent meta species index
$\square_x$	Axial component

Acronyms / Abbreviations

AE	Autoencoder
ANN	Artificial neural network
API	Application programming interface
BoP	Balance of plant
CCUS	Carbon capture, utilisation and storage
CFD	Computational fluid dynamics
`ChemZIP`	New ML-enhanced efficiently coupled aerochemical modelling methodology
CPU	Central processing unit
CSTR	Continuously stirred-tank reactor
DAC	Direct air capture
DACM	Dynamic adaptive chemistry method
DCZ	Dynamic cell zoning
DI	Direct integration (of the chemical kinetics)
EM	Electromagnetic
FGM	Flamelet generated manifold
FOM	Full-order model
GHG	Greenhouse gas emissions
GPU	Graphics processing unit
HHV	Higher heating value

IEA	International Energy Agency
ILDM	Intrinsic low-dimensional manifold
ISAT	*In situ* adaptive tabulation
LES	Large-eddy simulation
LCOE	Levelised cost of energy
LUT	lookup table
MAE	Mean absolute error
MAPE	Mean absolute percentage error
ML	Machine learning
`MPI`	Message passing interface
MSE	Mean square error
NMdAE	Normalised median absolute error
NN	Neural network
NRMSE	Normalised root mean square error
ODE	Ordinary differential equation
PCA	Principal component analysis
PFR	Plug-flow reactor
RANS	Reynolds-averaged Navier-Stokes
RDR	RotoDynamic Reactor (turbo-reactor)
RHS	Right-hand side
ROM	Reduced-order model
RPM	Revolutions per minute
SOEC	Solid oxide electrolyser cell
SWR	Shockwave reactor
TCI	Turbulence-chemistry interactions
TE/LE	Trailing/leading edge
TFLITE	TensorFlow Lite
TF	TensorFlow
URANS	Unsteady Reynolds-averaged Navier-Stokes

Part I

Aerothermodynamics

Chapter 1

Introduction

1.1 Decarbonising High-Temperature Industrial Heat

The cross-cutting problem of decarbonising industrial heat has been termed the "grand challenge" of thermal science and energy engineering by many researchers [1, 2] and is a critical hurdle in the development of a sustainable, net-zero modern society. Industrial heat decarbonisation, especially at high temperatures, has been plagued by technical and economic difficulties [3]. Figure 1.1 shows that about a third of global energy demand comes from the industrial sector, and most of this is from fossil-fuel-based heat.

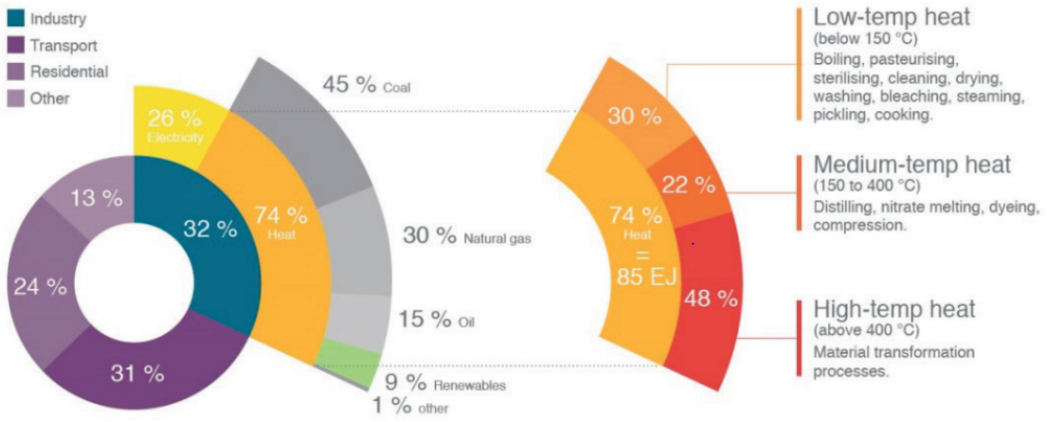


Figure 1.1: Breakdown of global energy demand by sector [4]

For low- to mid-temperature industrial heat ($T < 400\,^{\circ}\mathrm{C}$)—including drying, distillation, food, paper & pulp, boiling, etc.—relatively mature technologies are available that are technically and economically feasible. For example heat pumps, solar thermal, geothermal and electric heaters are potentially attractive solutions for this temperature range. However, these technologies are not applicable beyond 400 °C. For heat pumps, the low critical temperature and the thermal stability constraints of the refrigerants used in vapour compression cycles limit their applicability to temperatures below 150 °C. Even novel heat pumps based on air or steam cycles are restricted to a temperature level of

280 °C [5]. Theoretically, concentrated solar thermal could provide temperatures up to 3000 °C [6] but technologies in current operation are limited to ∼550 °C [2] and more importantly, the energy supply is intermittent. For geothermal energy, the main limitation is the depth of the drilling. Currently, in the oil industry, the maximum drilling depth technically feasible is 12 km. This limits the supply of geothermal heat to temperatures of up to 100 °C to 250 °C. However, if future developments allow for drilling depths of 100 km, temperatures of up to ∼1400 °C could theoretically be supplied [2].

It is clear that the so-called "grand challenge" is mainly focused on the decarbonisation of hard-to-abate, energy-intensive, *high-temperature* ($T > 400\,°C$) industrial heat, rather than lower temperature levels where more mature technologies exist. Figure 1.1 indicates that approximately 50% of industrial heat demand is high-temperature heat, equivalent to 10% of global primary energy consumption. This comes almost entirely from four vital industries: steel, cement, aluminium and chemicals/petrochemicals, where the vast majority of the thermal energy is provided by fossil-fuel-based combustion within industrial furnaces. These industries constitute some of the world's largest sectors with a market value close to $3 trillion in 2024, which could grow by 5.4% on average between 2024 and 2030. To put this into context, the aviation sector is valued at approximately $0.5 trillion with a growth rate 2 percentage points lower (3.4%).

High-temperature industrial sectors play a foundational role in modern society, producing building materials and high-value chemicals essential for synthesising critical raw materials such as plastics, synthetic fibres, rubbers, pharmaceutical products, and many more. However, to combat the climate emergency, greenhouse gas (GHG) emissions from these sectors must be eliminated. There are two components to high-temperature industrial GHG emissions: energy use and by-products of the chemical conversion process. Including both of these, steel, cement, aluminium and chemicals/petrochemical combined account for nearly 16% of global GHG emissions (see Fig. 1.2), which is around 8 times the emissions from aviation.

High-temperature industrial heat is evidently one of the most important but also most challenging sectors to decarbonise. A financially attractive, efficient and technically feasible solution is urgently needed.

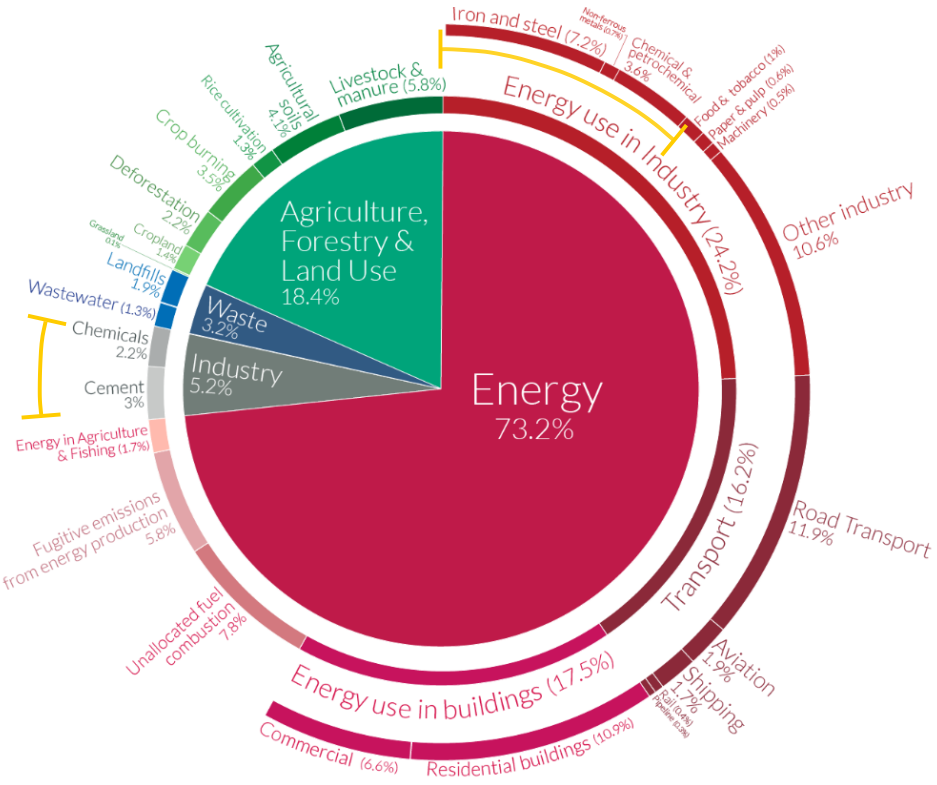


Figure 1.2: Breakdown of global greenhouse gas emissions by sector [7]

1.1.1 Direct Electrification of High-Temperature Industrial Heat with a New Turbomachine

Decarbonisation through *direct* electrification using zero-carbon electricity is potentially an energy efficient solution for high-temperature industrial processes. In contrast to power-to-fuel solutions, the energy conversion losses are mitigated. Currently, a critical barrier to the commercialisation of direct electrification solutions is the cost of electricity. The costs could be at least 4 to 15 times that of natural gas that is typically burnt within furnaces [8]. However, with a strong carbon tax and a carbon border adjustment mechanism, combined with a further decline in renewable electricity prices, Fig. 1.3 indicates that electrification could be cost-competitive when renewable electricity prices reach 30 – 50\$/MW h and the carbon tax is set to 100\$/t-CO_2. Figure 1.4 indicates that this is consistent with approximate projections for electricity and carbon prices over the next 40 years. Moreover, for industries such as the chemical industry, the cost of the feedstock far outweighs the cost of energy, as shown in Fig. 1.5. Therefore, one way to counteract high electricity

prices is through relatively small improvements in the efficiency of the reaction (*i.e.*, increasing the mass of high-value products produced for a given mass of feedstock).

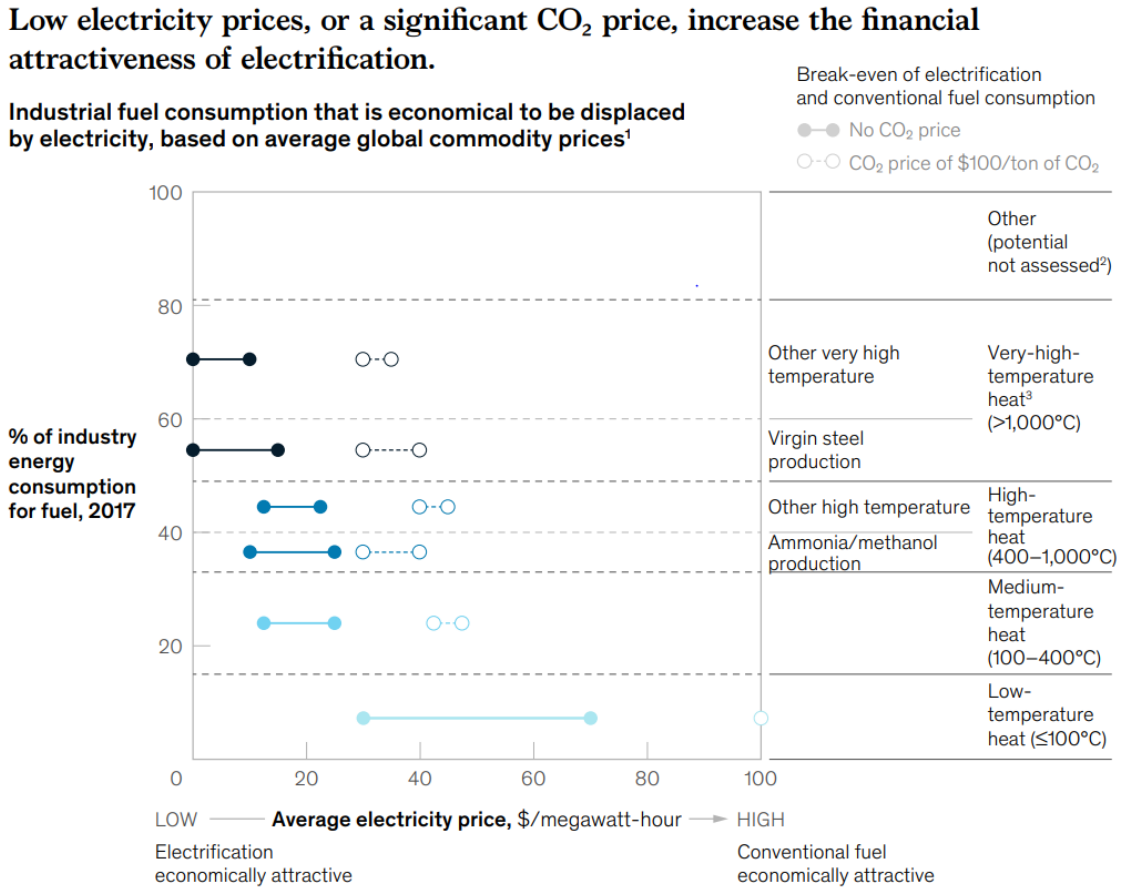


Figure 1.3: Electricity price and carbon tax necessary for industrial electrification of heat to become cost competitive [9]

The revolutionary solution to electrify industry introduced in this thesis is a new turbomachinery concept called the RotoDynamic Reactor (RDR) or turbo-reactor, as shown in Fig. 1.6. This concept replaces the ineffective and inefficient surface heat exchange energy transfer mechanism (within a furnace) with direct mechanical energy transfer through aerodynamic action within rotating blade rows. A fundamentally new design principle of the RDR is that all the mechanical energy imparted into the fluid is dissipated into internal energy rather than being converted into pressure of the fluid. By leveraging a renewably-powered electric motor as the mechanical driver, it is possible to electrify and decarbonise a broad range of high-temperature industrial processes.

What distinguishes the turbo-reactor from alternative pathways to decarbonise high-

temperature industrial heat (*e.g.,* electric heaters, hydrogen, etc.) is its significantly higher power density, enhanced process controllability, efficiency & operability, and extended machine lifetime. This is possible only by departing from surface-heat-exchange-based technology, which forms the foundation of most other low-carbon pathways. These features of the RDR lead to a reduction in costs compared to conventional technology, rendering it cost-competitive even under relatively high electricity prices and lower carbon prices.

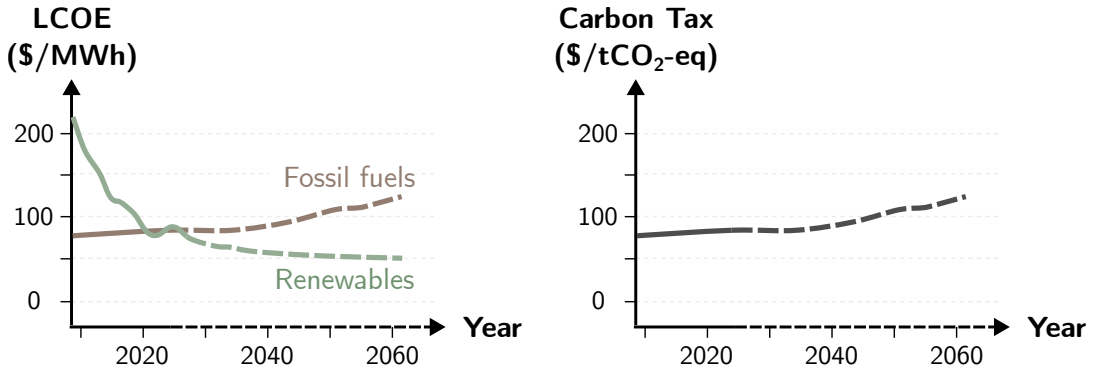


Figure 1.4: Schematic illustration of historical and 40-years projections for the levelised cost of energy (LOCE) including storage and carbon tax (based on data from [10, 11])

This thesis concentrates primarily on high-temperature endothermic chemical reaction applications of the RDR, as they present a new set of unique challenges and opportunities for turbomachines, unlocking an entirely new design space. In these applications, for the first time in turbomachinery, the working fluid chemically reacts and transforms (or decomposes) along the majority of the bladed gas path of the turbomachine. Therefore, interactions between the aerothermal flowfield and chemical kinetics must be taken into account to accurately predict the distribution of the primary product yield and coke formation rates. Ultimately, along with power input and power density requirements, yield and coking rates are the most important metrics of the process. As eluded to in Fig. 1.5, a promising vehicle to increase the cost-competitiveness of electrification is reducing feedstock consumption. Despite this, the current RDR design ecosystem does not account for aerothermochemical interactions, since incorporating chemistry into computational fluid dynamic (CFD) modelling is computationally intractable for detailed (but accurate) kinetic models. This problem is most acute for heavy high-carbon-number feedstocks, such as naphtha, for which the number of reaction pathways grows dramatically. Therefore,

neither the aerodynamics nor the reaction performance can be accurately predicted. While one-dimensional or homogeneous formulations are reasonable approximations for chemistry in furnace tubes [12], the same cannot be said for the turbo-reactor due to the complexity of the aerodynamic flowfield. For numerical combustion, there are established acceleration strategies to bring down the cost of computations, but these are either not sufficiently fast or not applicable to modelling reactive flows in a turbomachinery environment.

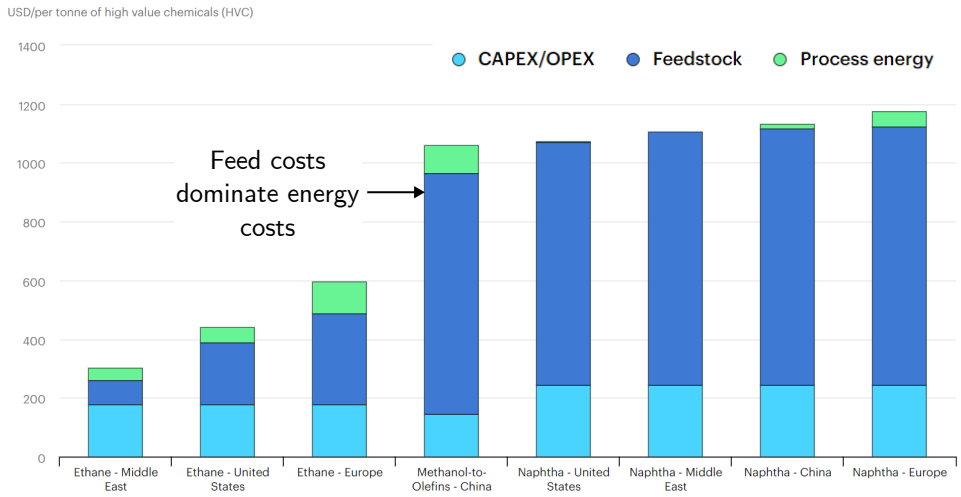


Figure 1.5: Cost breakdown for high-value chemicals production [13]

A crucial contribution of this thesis is developing a multi-fidelity machine-learning-assisted efficiently coupled aerochemical modelling methodology, called `ChemZIP`, to accelerate chemically reacting flow numerical simulations by two to three orders of magnitude. This is achieved while accurately capturing the response dynamics of a *detailed* (and realistic) kinetic model and maintaining a high-fidelity 3-D viscous aerothermal fluid model. Crucially, in contrast with other approaches, training times can be as low as a few hours and (2) the model does not impose additional time-step restrictions on the basic flow solver. This new methodology is an automated approach that leverages machine learning to uncover preexisting low-dimensional latent spaces hidden within the governing physics. As a result, aerothermochemical optimisation of the bladed path can now be performed early in the design cycle to improve the chemical reaction efficiency. Ultimately, this will enable cost-competitive electrification.

Throughout this thesis, the main endothermic chemical reaction focused on is the pyrolysis reaction of steam cracking. This is because, in the chemical/petrochemical



Figure 1.6: A turbo-reactor render [14] and some example industries that it could decarbonise

industry, steam cracking and ammonia production stand out as the main source of carbon dioxide (CO_2) emissions, accounting for 72% of total emissions [15]. For steam cracking, the primary contribution (90%) comes from the combustion of fossil fuels in steam cracking furnaces to drive the endothermic reaction [16].

However, it should be noted that the same fundamental turbomachinery elemental stage design is being applied to a broad range of other endothermic reactions (see Sec. 4.2 in Chapter 4), as well as for non-reacting high-enthalpy gas heating in the steel, cement and glass industries, and for sensible heat storage, methane pyrolysis and aqueous direct air capture technologies, as illustrated in Fig. 1.6.

1.2 Research Questions and Thesis Structure

The structure of this thesis is intended to reflect the inherently interdisciplinary nature of this engineering problem, which features strong interactions between aerothermodynamics and chemical kinetics. As such, this thesis is divided into two distinct parts. The first

part develops an understanding of the aerothermal flow physics of the turbo-reactor using decoupled chemical kinetic calculations, while the second part is centred around capturing aerothermochemical interactions at a low computational overhead. Therefore, it follows that this thesis will contain two condensed literature reviews, one in Chapter 2 (industrial heating and aerothermodynamics) and the other in Chapter 6 (chemistry + aerothermochemistry). Note that often in this thesis, "aerothermodynamics" is shortened to "aerodynamics", and "aerothermochemistry" is shortened to "aerochemistry".

This thesis is structured as follows. Chapter 2 introduces alternative pathways to decarbonise high-temperature chemical processes and discusses their challenges and limitations. Subsequently, the new turbo-reactor concept is introduced. The next two chapters concentrate on understanding the aerodynamics of the axial turbo-reactor with decoupled and highly simplified chemical kinetics. These chapters address the following research questions:

1. *Is the turbo-reactor an efficient and effective solution to impart a large amount of mechanical energy into a working fluid over short length and time scales?*

   After introducing the new turbo-reactor concept, Chapter 2 details a series of numerical simulations that demonstrate the feasibility and robustness of using a novel regenerative turbomachinery architecture to decarbonise high-temperature processes for the first time.

2. *What are the physical mechanisms responsible for the conversion of mechanical energy into internal energy?*

   Chapter 3 first explains how the RDR can be a "chemically tuned" energy conversion device, hopefully correcting the common misconception that it is just a "bad compressor". Next, high-fidelity numerical investigations reveal the breakdown of energy conversion mechanisms. This will help improve the temperature profile for more reaction-efficient future design iterations.

3. *Is the elemental stage design concept robust & flexible and can a universal stage design philosophy be used for a wide range of working fluids?*

   Chapter 4 investigates the impact of the interstage distance on the stage aerodynamics. The minimum vaneless space length possible while still maintaining well-conditioned flow sets the lowest allowable reaction time-scale. Given the flexibility of the aerodynamic design, the advantages of redistributing the stages in a multi-stage environment to improve reaction performance are demonstrated. Subsequently, the impact of various global stage parameters and operating conditions on the reaction performance are investigated.

   Finally, a comprehensive analysis of aerothermal performance across a broad spectrum of feeds, chemical reaction stages and Mach number regimes is used to determine the feasibility of a universal stage design philosophy.

Chapter 5 serves as a bridge between the aerothermodynamics-focused ("Part 1") and chemical- & aerothermochemical-focused ("Part 2") chapters. This chapter highlights the unique level of complexity of the flow physics and the new requirements and tools necessary to tailor the design ecosystem to the turbo-reactor. The final three research chapters are related to chemical and aerothermochemical modelling in the turbo-reactor and intend to address the final research question:

4. *Can a multi-fidelity, data-driven methodology be used to accelerate the modelling of aerochemical interactions in the turbo-reactor with sufficiently high accuracy that they can feature routinely in parametric design optimisation and exploration?*

   Chapter 6 starts by emphasising why aerochemical interactions are important in the turbo-reactor. After this, the relevant literature on techniques to accelerate chemical kinetic modelling is reviewed, and the limitations of these methods are identified. To overcome these restrictions, Chapter 7 presents a new data-driven methodology. Chapter 8 demonstrates the feasibility of the `ChemZIP` methodology through a sequence of comprehensive performance benchmarking verification studies.

Finally, Chapter 9 summarises and discusses the main findings of this thesis and suggests interesting directions for future work.

Chapter 2

A New Turbomachinery Concept to Decarbonise High-Temperature Processes

2.1 A Review of Concepts to Decarbonise Chemical Processes

Converting chemical transformation processes (*e.g.,* steam cracking, steam methane reformation, methane pyrolysis, methanation, etc.) into fully sustainable, carbon-neutral industries requires tackling three emission and non-renewable sources: process heat, chemical feedstocks, and by-products of the chemical conversion process. The latter is especially important for ammonia production, where 60% of the GHG emissions are released from CO_2 as a by-product of the chemical reaction [17]. There are various solutions to this, such as changing the chemical reaction iteself—for example, to an electrochemical reaction [17]—to circumvent CO_2 as a by-product, or by using carbon capture, utilisation and storage (CCUS). To address the second non-renewable source, the chemical feedstock, there are three important pathways [18] (1) biomass-derived feeds (2) synthetic feeds derived from combining CO_2 from direct air capture (DAC) with hydrogen (H_2) produced from water electrolysis [19] and (3) increasing the circularity of chemicals by manufacturing feedstocks from thermochemically recycled plastic waste [20]. It is emphasised that decarbonising the feedstock extraction and production process is vital, since not only is it unsustainable, but it can also contribute to more than 50% of the lifecycle GHG emissions from high-value chemical production [21].

This literature review concentrates on the first of the three emission and non-renewable sources: process heat. The following sections first present a brief review of conventional furnace technology for endothermic processes followed by a comprehensive critical analysis of state-of-the-art solutions to decarbonise high-temperature heat in the chemical industry,

underlining the challenges and limitations of each technology. Subsequently, a new turbomachinery concept is introduced to decarbonise process heat and address many of the limitations of alternative pathways.

Before moving forward, it is helpful to formally define the following metrics in the context of chemical processes:

- *Reaction product yield (Y_k)* = the percentage mass of the k^{th} chemical species relative to the total gas mixture mass (*i.e.,* the k^{th} species mass fraction)
- *Reaction/feed (mass) conversion (C_{feed})* = the percentage mass of feedstock converted into other products ($C_{\text{feed}} = 1 - Y_{\text{feed}}$)
- *Reaction selectivity* = the ratio of a high-value primary product yields to the fraction of feedstock converted (*i.e.,* Y_k/C_{feed})
- *Reaction efficiency* = percentage of feedstock mass that ends up as desired primary product species ($\sum_k^{N_{\text{desired}}} Y_k$)

For chemical processes that use steam as an inert diluent, yields are often presented as *dry* mass fractions. This means that if the gravimetric fraction of steam at the inlet is 50%, a true yield of 40% translates to a dry yield of 40% / 50% = 80%. The convention of using dry yield will be adopted throughout.

2.1.1 Conventional Hydrocarbon Cracking Process

Before examining concepts to decarbonise general endothermic chemical reactions, it is beneficial to deconstruct an example application of conventional fossil-fuel-based chemical reactions: steam cracking of hydrocarbon feedstocks [22, 23]. Steam cracking, also known as pyrolysis, is an endothermic thermochemical decomposition reaction that breaks heavy, *low-value* long-chain hydrocarbon molecules such as butane (C_4H_{10}) into low-molecular-weight *high-value* olefins such as ethylene (C_2H_4). This is typically a non-catalytic process; however, many other chemical processes such as ammonia cracking require catalysts to increase reactivity.

Figure 2.1 highlights several key components of a typical steam cracking plant. The convection section recuperates heat not absorbed in the radiant section to preheat the unreacting feedstock (*e.g.*, naphtha) and dilution steam. The feedstock is then mixed with the steam to reduce the hydrocarbon partial pressure, which reduces the coke formation rate. In the gas-fired radiant section, endothermic reactions proceed rapidly and olefins are produced. Once the reaction is complete, the products are quenched, compressed and separated (see Fig. 2.1).

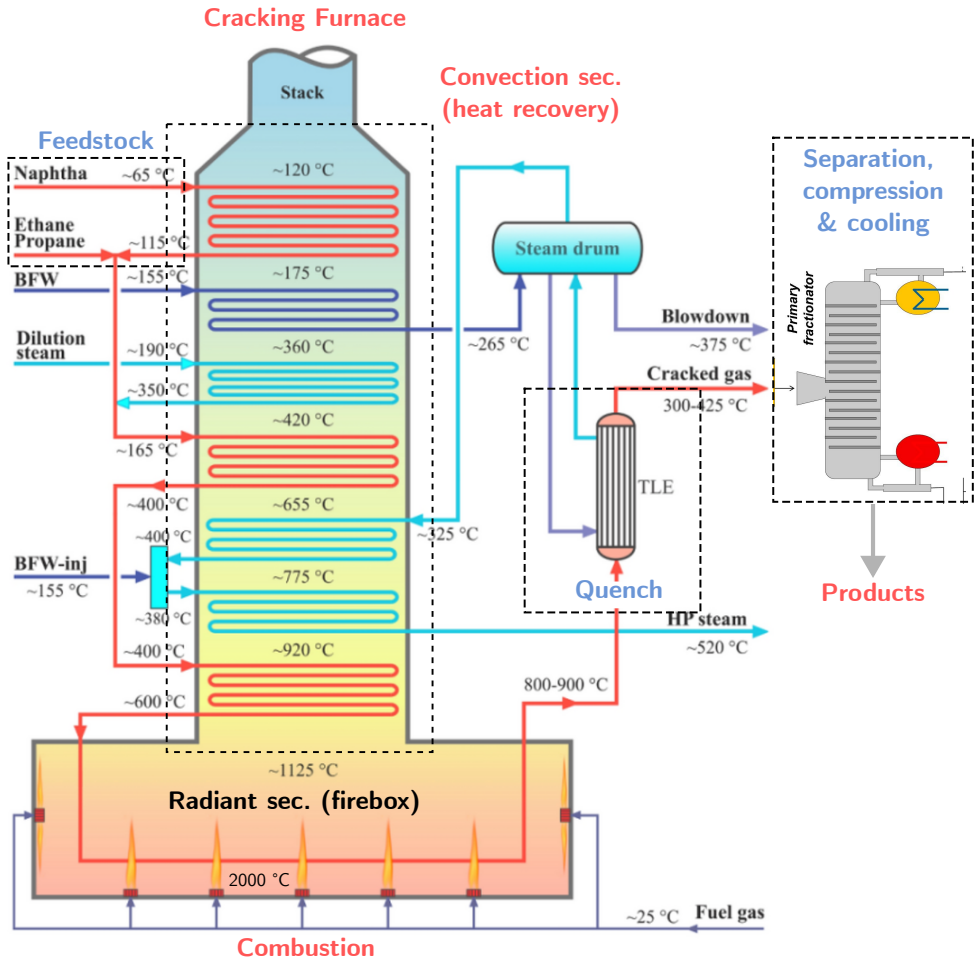


Figure 2.1: Conventional steam cracking plant schematic (adapted from [24, 25])

Fundamentally, these furnaces are surface heat exchangers that require very large temperature differences to drive energy transfer (see Fig. 2.1). As a result of the surface heat exchange mechanism, conventional technologies suffer from low power density, short coil lifetime, restricted reaction efficiency and heavy carbon deposition on the walls.

Coke Formation and Deposition

For hydrocarbon-based chemical reaction processes, solid carbon deposition on metal surfaces—*i.e.,* coking, as shown in Fig. 2.2—is a critical challenge. Solid carbon is formed directly on the metal surfaces through heterogeneous catalytic and pyrolytic gas-to-surface phase reactions [28–30]. Coking leads to three key challenges: (1) it severely degrades

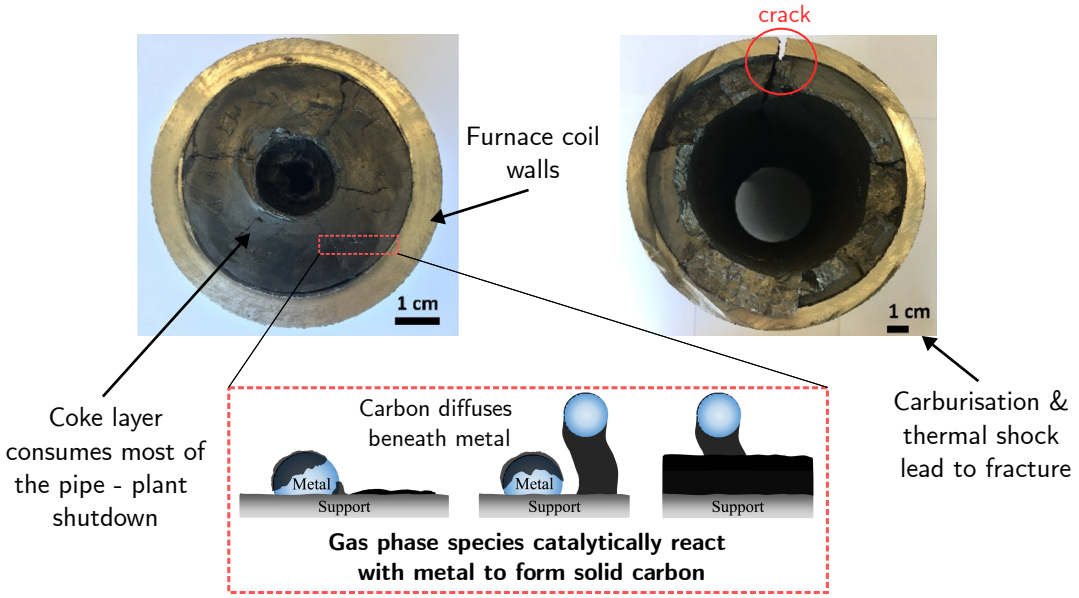


Figure 2.2: Coke accumulation within a conventional ethylene furnace coil (Refs. [26, 27])

the service life of the furnace coils due to carburisation [31] (see Fig. 2.2) (2) it increases heat transfer resistance through the walls [26] and (3) it reduces plant availability due to periodic decoking cycles. The service life of the coils can be as low as 4 to 6 years [32] and periodic replacement can cost 3 to 7% of the entire plant [33]. Furthermore, decoking [34] must occur as often as once a month for 2 to 3 days at a time [30, 35], leading to a large loss of annual production capacity, as well as life-degrading thermal cycles. Consequently, in many applications, mitigating coking is equally as important to the designer as optimising the primary product yields. However, in conventional surface-heat-exchange-based furnaces, there is limited scope to minimise coking. This is because coke formation is driven by high gas-coke interface temperatures, which cannot be significantly reduced while maintaining the required heat transfer rate for the reaction.

Now that conventional furnace technology has been briefly reviewed, the following sections explore alternative technologies to decarbonise these high-temperature

endothermic chemical processes.

2.1.2 Power-to-X: Carbon-Neutral Fuels

Power-to-X concepts convert renewable electricity into carbon-neutral fuels such as hydrogen or synthetic methane, which can then be combusted in furnaces. These technologies are popular because existing pipeline infrastructure and/or assets within high-temperature industries—radiant furnaces, boilers, preheaters, etc.—can remain in use with either no changes or relatively minor modifications to the combustor.

A critical performance metric for these technologies is the energy conversion efficiency, which is defined as follows:

$$\eta_{\text{conversion}} = \frac{\text{HHV}_{\text{fuel}}}{Q_{\text{in}} + P_{\text{in}} + E_{\text{chem.}} + P_{\text{trans.}}} \tag{2.1}$$

where HHV_{fuel} is the higher heating value of the fuel produced (*e.g.,* synthetic methane); $E_{\text{chem.}}$ is the heating value of the reactants, Q_{in} & P_{in} are the heat and electrical power inputs required for the conversion process including additional energy required for auxiliary supporting subsystems (*e.g.,* inverters, rectifiers, cooling, purification, etc.), *i.e.,* the balance of plant (BoP); and $P_{\text{trans.}}$ is the power loss during storage, transport, & distribution of the fuel. This metric represents the energy content of the fuel produced per unit of energy consumed to create it.

Synthetic Methane

Synthetic methane (CH_4) is produced by combining CO_2 from a DAC system with H_2 from water electrolysis [36]. The energy supply for these processes is zero-carbon electricity. The exothermic Sabatier methanation reaction is used to convert the CO_2 and H_2 reactants into synthetic methane, as shown below:

$$\text{CO}_2 + 4\,\text{H}_2 \xrightarrow[\text{pressure + catalyst}]{300\,^\circ\text{C}} \text{CH}_4 + 2\,\text{H}_2\text{O} \quad (\Delta H_f^0 = -164\,\text{kJ mol}^{-1}) \tag{2.2}$$

where ΔH_f^0 is the standard enthalpy of reaction. Alternatively, a review by Biswas *et al.* [36] identifies *in situ* methanation within a solid oxide electrolysis cell (SOEC) as an

ultra-efficient pathway to produce synthetic methane. Higher efficiency is achieved by integrating H_2 production and methanation into a single SOEC package/unit.

The two main drawbacks of synthetic methane are high capital & operating costs and low conversion efficiency. Using the techno-economic analysis presented by Kiani *et al.* [37], the conversion efficiency could be as low as $\eta_{\text{conversion}} \approx 45\%$ (based on Eq. 2.1), including transportation and furnace heat losses (5%). The *in situ* methanation SOEC discussed above could yield a 20% improvement in conversion efficiency, but this concept is in the early research phase [36]. Furthermore, in terms of levelised cost of energy (LCOE), Kiani *et al.* estimate that synthetic methane could cost 15 times more than natural gas (*i.e.,* 115\$/MW h). Almost 70% of this cost is associated with the cost of electricity used for electrolysis.

Although e-methane is a drop-in replacement for conventional natural gas, its additional costs and conversion efficiency losses limit its commercial value. Generally, synthetic fuels are better suited for aviation and shipping, where space is premium.

Biomethane

Biomethane production shares similarities with synthetic methane production. However, in this case, biogas (CH_4 + CO_2) is first produced from anaerobic digestion of organic matter by micro-organisms in an oxygen-free environment [38], as shown in Fig. 2.3. The feedstock for biogas production can be municipal solid waste, landfill gas, animal manure, and crop waste residues [39]. The waste by-product of the reaction is digestate, which can be used as an organic fertiliser [39].

Next, the biogas is purified and upgraded to convert it into biomethane [38], where the CO_2 fraction is either removed and released (see Fig. 2.3), or converted to more methane by combining it with green hydrogen through a methanation reaction (see Eq. 2.2). This separation process is substantially less energy intensive than carbon capture from the flue gas of a power plant because the concentration of CO_2 is 10 to 20 times higher. The characteristics of biomethane are almost identical to those of methane, and so it can be used as a drop-in replacement.

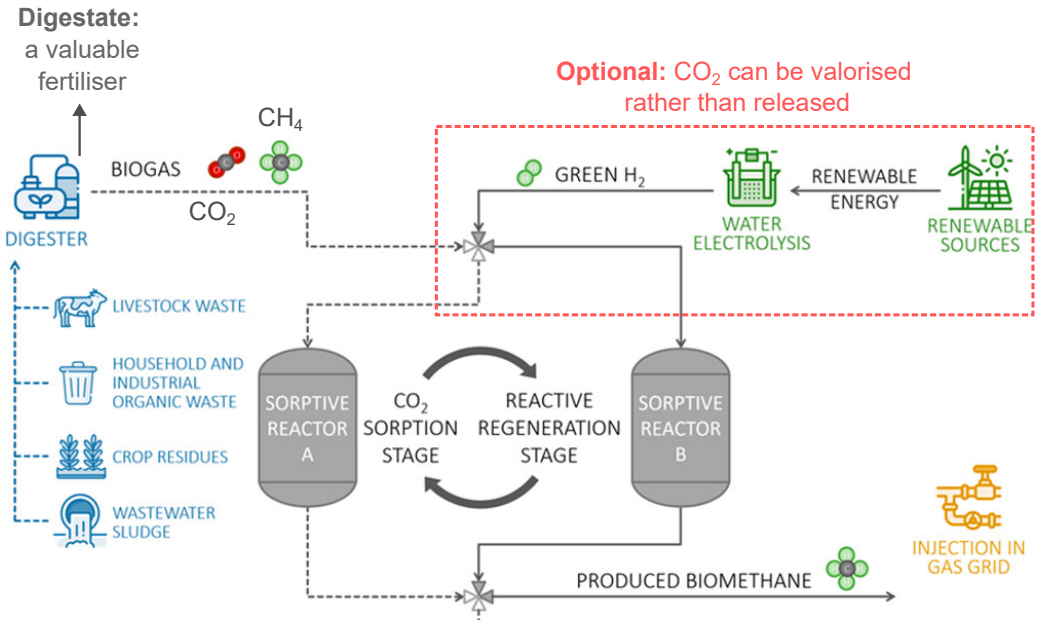


Figure 2.3: Biomethane production process [40]

The main advantage of biomethane is that its cost could be about half that of synthetic methane at around 50 – 90 \$/MW h [41, 42]. This is because (1) many biomethane feeds are almost free (*e.g.,* landfill gas), (2) biomethane avoids the expensive DAC system and (3) the main cost, *i.e.,* the biodigester, operates at low temperatures.

However, the International Energy Agency (IEA) estimated a maximum annual technical potential for biomethane of ∼8100 TW h [41] in 2022, which is only enough to meet about half of the heat requirements of high-temperature industrial processes.

Green Hydrogen

Combusting hydrogen directly within high-temperature industrial furnaces bypasses the need for a DAC system that was required for synthetic methane production (see Sec. 2.1.2). This enables overall conversion efficiencies of $\eta_{\text{conversion}} \approx 60\%$ [37], which is significantly better than for synthetic methane production.

Generally, the financial viability of hydrogen as a carbon-neutral fuel (and also with the e-fuel in Sec. 2.1.2) is highly dependent on the conversion efficiency of the electrolyser due to the high prices of electricity relative to the prices of gas. The electrolyser system efficiency can vary substantially from as low as 50% up to a maximum of 80% [43].

Recently, however, Hodges *et al.* [44] developed a novel low-temperature capillary-fed electrolyser at lab-scale (see Fig. 2.4) that produces a conversion efficiency of more than 90% (including BoP) by eliminating the formation of performance-degrading bubbles at the electrode-electrolyte interface.

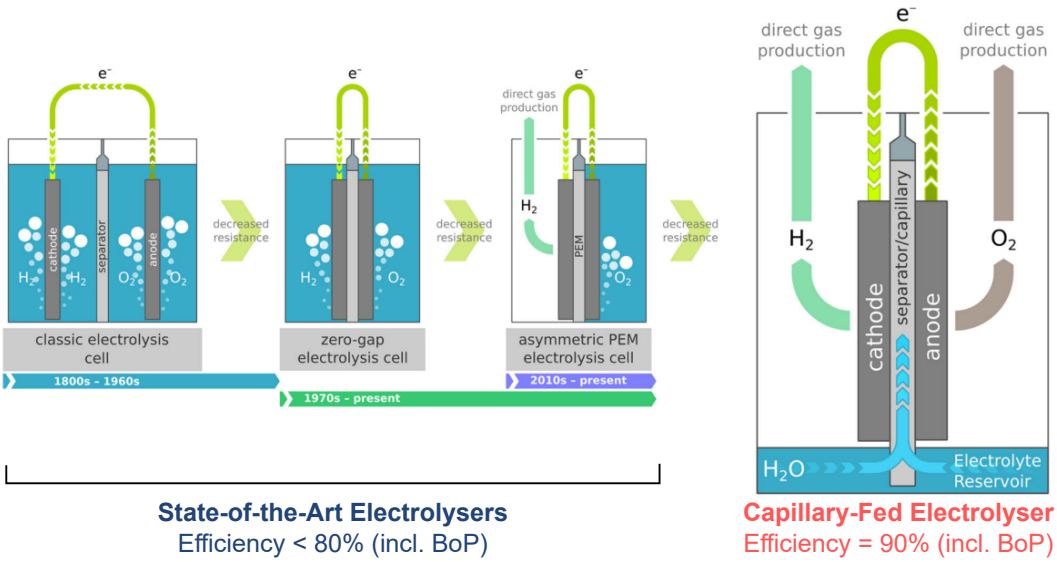


Figure 2.4: Green hydrogen production via water electrolysis [44]

There may be additional energy loss, between the 6% and 8% [43], due to hydrogen compression and transport (*e.g.,* along a 100 bar pipelines). Therefore, decentralised hydrogen production on site is expected to be more cost-effective, eliminating transportation losses and costs [45]. This is particularly true for plants in the chemical industry. If, instead, a centralised hydrogen hub is used due to economies of scale benefits, a comprehensive study by DeSantis [46] indicates that the cost of hydrogen transport across a pipeline network is 8 times cheaper than the transmission of electric power. This is a notable advantage of hydrogen for the decarbonisation of high-temperature industries in comparison with direct electrification. However, H_2 pipeline infrastructure is virtually non-existent and may take decades to build at scale if natural gas pipelines are unsuitable to be repurposed.

Despite the advantage in conversion efficiency compared to synthetic methane, Giacomazzi *et al.* [47] points out that the challenges associated with scaling-up hydrogen combustor technology have not been fully solved. These include flashback risks, thermoacoustic instability and NO_x emission control. However, many of the hurdles to hydrogen combustion physics can be resolved with novel combustors such as micromix burners

[48] and two-stage combustors, which exhibit low NO_x emissions and enhanced flame stability. Unfortunately, these are still in the early stages of development. Moreover, compared to synthetic methane and biomethane, substantial modifications are required to the existing gas pipeline network.

Other

Throughout the discussion above, we have ignored blue hydrogen [21] and CCUS [49]. This is because the ultimate goal is complete sustainability throughout the lifecycle, and both these technologies use non-renewable natural gas as energy sources.

2.1.3 Electric Furnaces

Direct electrification with electrically heated furnaces eliminates many of the conversion efficiencies that drive up the cost of power-to-X technologies. The main energy conversion losses for electric furnaces occur in the electrical transmission lines. Over a distance of 500 miles, losses can exceed 6% [46, 50]. The magnitude of this loss is highly dependent on the proximity of the renewable energy source (and storage) to industrial plants (*i.e.,* is power generation centralised or decentralised?). Once electricity is received at the plant, the thermal efficiency of its conversion into heat can be close to 100%.

Despite a large body of research on electric heaters, as reviewed by Zheng *et al.* [51], they have yet to be commercially implemented on a large scale for high-temperature, high-capacity industrial processes. However, many chemical producers are currently developing pilot e-furnaces [52, 53].

Joule Heating

Joule heating [54–56], also known as ohmic heating, operates by passing a current through an electrical conductor. Charged particles collide with atoms in the conductor, causing atomic vibrations that convert electric energy into thermal energy. This physical process can theoretically convert 100% electrical energy into heat [57]. However, the efficiency depends on the material properties. Metals, such as FeCrAl [58], or ceramics, such as SiSiC [54], provide a good balance between temperature resilience and electrical conductivity.

The electrically heated elements can be located within the walls of the furnace tubes, as shown in Fig. 2.5, or as structured inserts within the flow path of the tube.

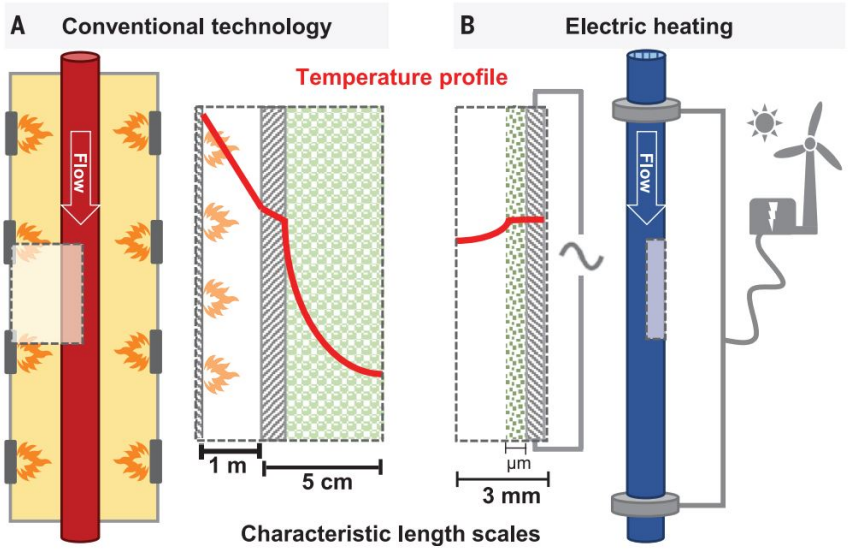


Figure 2.5: Joule heated electrified steam methane reforming reaction for H_2 production [56]

Joule-heated electric furnaces provide considerable advantages compared to combustion-fired furnaces. Most importantly, the tube walls or support structures inside are directly electrically heated. This reduces the maximum temperature of the system and minimises temperature gradients [59], which can substantially improve the selectiveness and efficiency of the reaction. Perhaps the most powerful application of joule heating in the chemical industry is for catalytic processes. In novel designs, catalytically coated porous metallic foam-based support structures can be electrically heated directly [60]. This provides intimate contact between the heated substrate, the catalytic washcoat and the gas. Therefore, heat transfer is no longer the limiting time scale and thus the reaction rate and performance can be improved [59].

Electromagnetic Heating

Electromagnetic (EM) heating can be found in many forms: induction, microwave, etc., but the common fundamental principle among them all is that many materials can absorb EM waves and convert them into heat [61]. For induction heating, a coil wrapped around the furnace tube produces an alternating magnetic field that induces eddy currents within the susceptor material (*e.g.,* the furnace tube, catalyst pellets or foam insert). These eddy

currents generate thermal energy, but are induced only in specific materials with magnetic properties. Figure 2.6 shows a significant advantage of induction heating: it can induce heat directly within catalyst pellets that are in intimate contact with the reacting fluid.

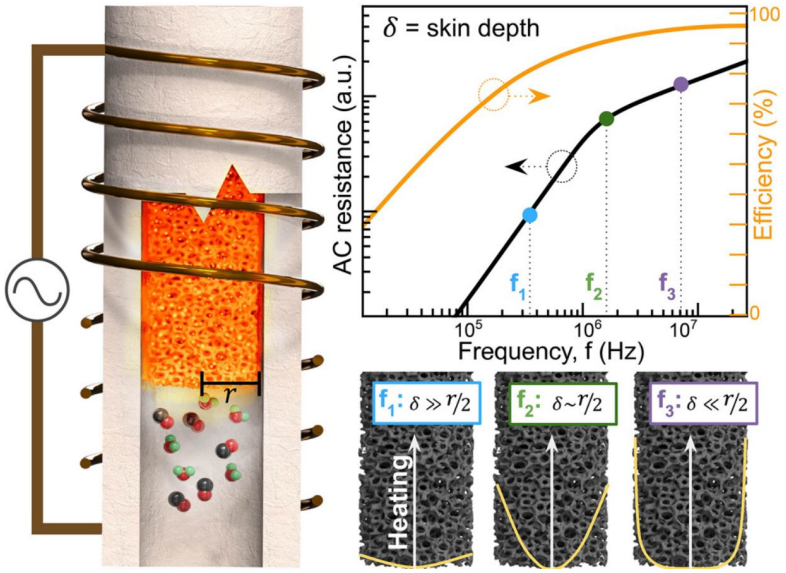


Figure 2.6: Magnetic induction heating for thermochemical processes using metamaterial-based foams [62]

Alternatively, microwaves generated in a circuit decoupled from the furnace can be absorbed by a microwave-absorbing material. This cannot be a metal as it reflects microwave radiation. Instead, dielectric materials such as SiC are used. More interestingly, for feedstocks mixed with steam, the gas mixture can absorb microwaves and be heated directly [61, 63]. This could potentially result in a more compact reactor that exhibits a homogeneous, and hence high-performance, chemical reaction.

Compared to Joule heating, EM heating has the advantage that there is no need to insert electrical elements into the furnace. The electrical circuit can be decoupled. This can also allow precise temperature control and targeted heating [61]. However, for all EM routes, the energy efficiency is inevitably lower than for Joule heating, which can be around 95% [64]. For induction heating, the thermal efficiency at laboratory scale is reported to be not much more than 20% [65] and microwave heating at best 75% [61, 66]. At best, Lin *et al.* [62] suggest that a novel metamaterial reactor powered using high-frequency magnetic induction (see Fig. 2.6) could achieve an efficiency of 90% once scaled to a commercial level. However, there are limited susceptor materials where the

Curie temperature—the temperature above which magnetism is lost—exceeds the reaction temperature (500 – 1000 °C) [59]. Ceramics such as SiCSiC will likely have to be used [62].

Non-Thermal Plasma Heating

Plasma reactors fall into two categories: Thermal and non-thermal plasma. Non-thermal plasma can be generated by applying a voltage that is large enough to partially ionise the gas (see Fig. 2.7), unlocking the free movement of electrons [61]. A remarkable property of this type of plasma is that the effective *electron temperature* (∼10,000 K) can become much higher than the bulk gas temperature (∼500 K) [67] by gaining energy from the applied electric field. Collisions between molecules and electrons can excite chemical reactions at gas temperatures significantly milder than conventionally possible [68]. Non-thermal plasma reactors have not been developed commercially, but Mehta *et al.* [69] reviews numerous developments by academic researchers.

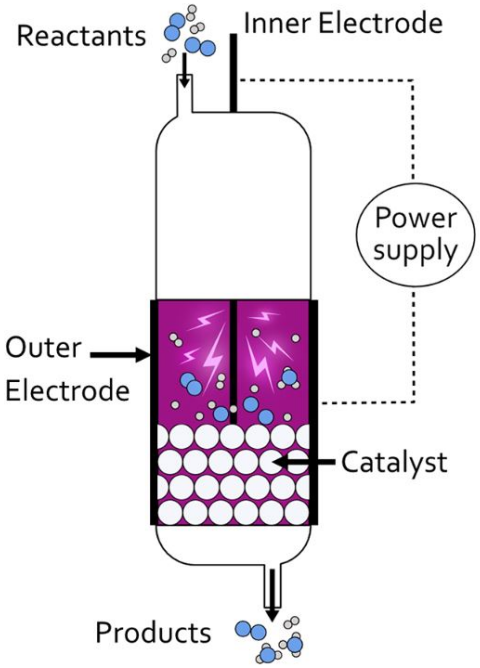


Figure 2.7: Schematic of a dielectric discharge non-thermal plasma chemical reactor [69]

2.1.4 Electrochemical

Although our focus is on decarbonising high-temperature industrial heat, it is important to assess whether thermally driven chemical processes are even necessary in the first place. In other words, is heat necessary? For chemical/petrochemical reaction processes, an efficient and effective method of electro-decarbonisation is the use of an electrochemical process [70]. Rather than driving the chemical reaction through a temperature gradient, a potential difference is used as the driving force. This permits a significantly lower operating temperature ($T < 100\,^{\circ}\mathrm{C}$), providing advantages in terms of material selection, as well as almost eliminating unwanted side reactions and coke deposition (which are induced at high temperatures) [71].

Reports from Ding *et al.* [71] indicate that a highly selective (nearly 100%) electrochemical ethane cracking reaction is possible at a 35% lower energy consumption than conventional cracking. This suggests that electrochemical pathways could be important industry disruptors for achieving electro-decarbonisation. However, electrocatalytic materials that are both highly selective and have long lifetimes (>10 years) under large dynamic loads have not yet been conclusively demonstrated at scale [70]. Moreover, for other novel electrochemical processes such as ammonia synthesis (for fertiliser production), there have been challenges with low reaction conversion, reaction selectivity and faradaic efficiency [72].

2.1.5 Outlook

We have reviewed eight low-carbon technologies that could enable decarbonisation of high-temperature industrial heat, with a focus on the chemical industry. This outlook section will summarise the most efficient and effective pathways to decarbonise industrial heat. Some of the important comparative metrics are found in Table 2.1. However, this data should be taken with care. There is still a high level of uncertainty in cost estimates, and these can vary dramatically from region to region depending on available technology, natural resources and geopolitical factors.

For most technologies, perhaps with the exception of biomethane, large-scale adoption is critically dependent on low renewable electricity prices relative to fossil fuel energy

prices. With an effective carbon tax and falling renewable electricity prices (owing to economies of scale and zero marginal operating cost), electro-decarbonisation may be financially competitive in the near future [73].

Table 2.1: A summary of conversion efficiency and cost of energy for different methods to decarbonise high-temperature industrial heat [46, 73, 74]

Energy carrier	$\eta_{\text{conversion}}$ (%)	LCOE ($/MW h)	Transmission† ($/MW h)
Synthetic CH_4	45 – 65	115 – 130	2.0
Biomethane	85 – 95	50 – 190	2.0
H_2 production	60 – 80	110 – 150	2.5
Wind elec. + storage	85‡	40 – 115	20
Grid elec.	85‡	50 – 150	20
Natural gas	93‡	7 – 15	2.0

† assuming transmission over a distance of 500 miles
‡ the only energy loss is from transmission. Thermal furnace efficiency = 95%

With *Indirect* electrification through power-to-X technologies, biomethane cannot be scaled sufficiently to meet demand, and hydrogen is too costly (see Table 2.1). On the other hand, *direct* electrification (*i.e.,* joule heating, induction heating, etc.) is the cheapest and most efficient renewable pathways, and provides additional furnace compactness and reaction efficiency benefits as a result of improved heat transfer. Joule heating provides the highest theoretical efficiency (100%), lowest complexity and widest range of usable materials (independent of magnetic properties), so it is perhaps the most likely pathway for electric heaters. However, improvements are still required to minimise energy conversion losses, combat oxidation challenges, and demonstrate scalability and durability over extended periods. Moreover, the LOCE + transmission is still 5 to 15 times higher than natural gas (see Table 2.1).

Currently, grid transmission costs account for up to a third of the total energy cost. Across a distance of 500 miles, electrical transmission line losses (see Table 2.1) are around 6.5% and the cost of transmission is 20$/MW h [46, 50]. This is compared to 1.0 – 1.5% losses and 2 – 2.5$/MW h transmission cost for gas pipelines [46, 50] due to the significantly higher energy-carrying capacity. However, recent estimates from Shafiee & Schrag [75] suggest that the cost of hydrogen storage and distribution could be as high as 50 $/MW h, but the validity of this is uncertain.

One potential mechanism to reduce conversion losses and transmission costs is to locate high-temperature industries near renewable-energy-producing sites. Ideally, these integrated energy and industry hubs would be situated in regions with cheap (LCOE $\leq$ 30\$/MW h) and abundant renewable energy located near the sea for easy access to shipping ports. Figure 2.8 shows that these include regions in Africa, the Middle East, Australia and South America. Since high electricity costs are the main barrier to direct electrification of industries, this could have a substantial impact.

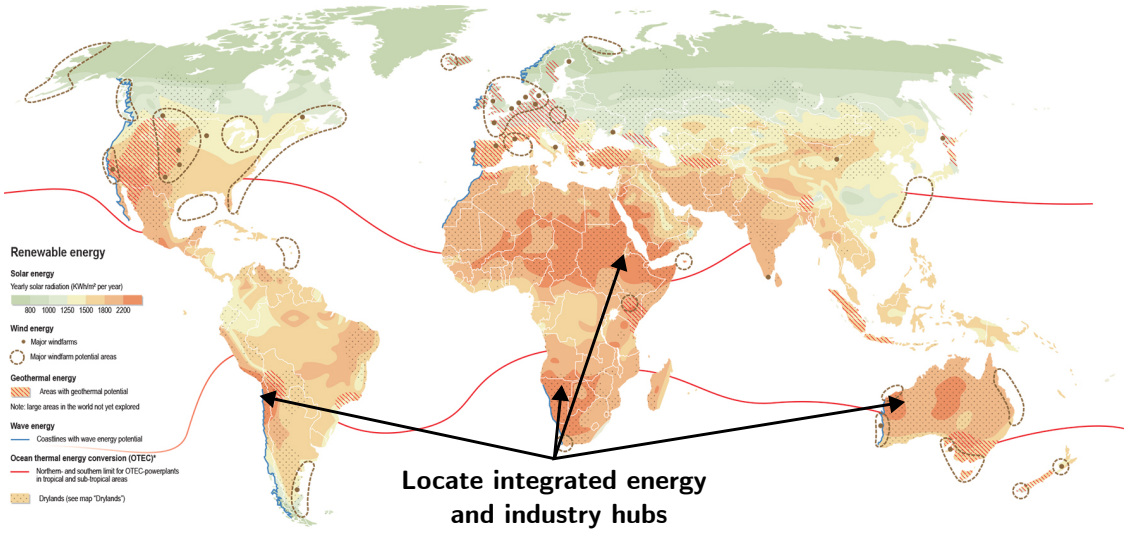


Figure 2.8: Global map of renewable energy resource potential [76]

Finally, from this discussion, it is clear that high electricity prices relative to gas prices are a key hurdle for most of the carbon-neutral heating pathways mentioned. Therefore, to increase the attractiveness of novel solutions, they should provide supplementary cost-reducing mechanisms such as reduced capital and maintenance costs and/or increased reaction efficiency and component lifetime. However, since most of the proposed carbon-neutral heating avenues—with the exception of plasma, electrochemical and microwave feed heating—are still surface-heat-exchange-based technologies, they are subject to the same limitations as conventional furnaces. The metal surface temperature must be hotter than the working fluid to provide heat transfer. This fundamental physical restriction prevents the step-change improvements in reaction efficiency, plant volume, and in-service life necessary to bring the cost of carbon-neutral solutions in line with conventional gas-fired furnaces.

2.2 A New Turbomachine for Clean Pyrolysis

2.2.1 Introduction

To address many of the limitations of the carbon-neutral pathways presented in the preceding section (Sec. 2.1), this thesis introduces a new electric-motor-driven turbomachine called the turbo-reactor. This machine can decarbonise high-temperature chemical processes, such as hydrocarbon cracking (also known as steam cracking or olefin production), through direct electrification. The turbo-reactor addresses most of the challenges of alternative pathways by replacing surface heat exchange with direct mechanical energy transfer through a rotating blade row. Therefore, the average metal temperature must be less than the working fluid temperature. This enables a step-change improvement in the quality of the process, providing the necessary cost-reducing mechanisms for financial viability.

Inspirations for RDR Concept

To meet the requirements of high-temperature industrial processes, the elemental stage and multi-stage environment of the RDR draws on several existing engineering and turbomachinery systems. For example, to achieve ultra-high specific enthalpy input, the RDR's bladed path shares similar features with both conventional supersonic impulse compressor designs developed in the 1950s [77] and unconventional supersonic turbo-compressors [78] developed in the early 2000s, as well as impulse steam turbines and turbines for pressure-gain combustion cycles [79].

In the early 1990s, Hertzberg *et al.* [80] developed a shockwave reactor (SWR) for pyrolysis [81–84] that mitigated the limitations of the surface heat exchange energy transfer mechanism in the thermal cracking zone. This enables shorter reaction times-scales. The SWR design principles has influenced the elemental stage construction of the turbo-reactor.

Figure 2.9 shows that Hertzberg's concept attempts to approximate volumetric heating of a chemical feedstock (*e.g.,* ethane) by first superheating an inert carrier gas (*i.e.,* steam) to a temperature of 1400 K, before injecting an un-reacted feedstock into the hot carrier gas. The two fluids supersonically mix-out to restore uniformity. Subsequently, a standing

normal shockwave (see Fig. 2.9) is setup to bring the gas mixture up to the reaction temperature, where the endothermic pyrolysis chemistry proceeds homogeneously in all spatial directions normal to the flow direction. Effectively, the heating processes and reaction process have been decoupled. In the reaction zone, this allows higher temperatures within a shorter reaction time-scale, directly translating to higher reaction selectivity (by up to 20%). Furthermore, because heat is not transferred through the walls, the casing can be cooled to eliminate coking.

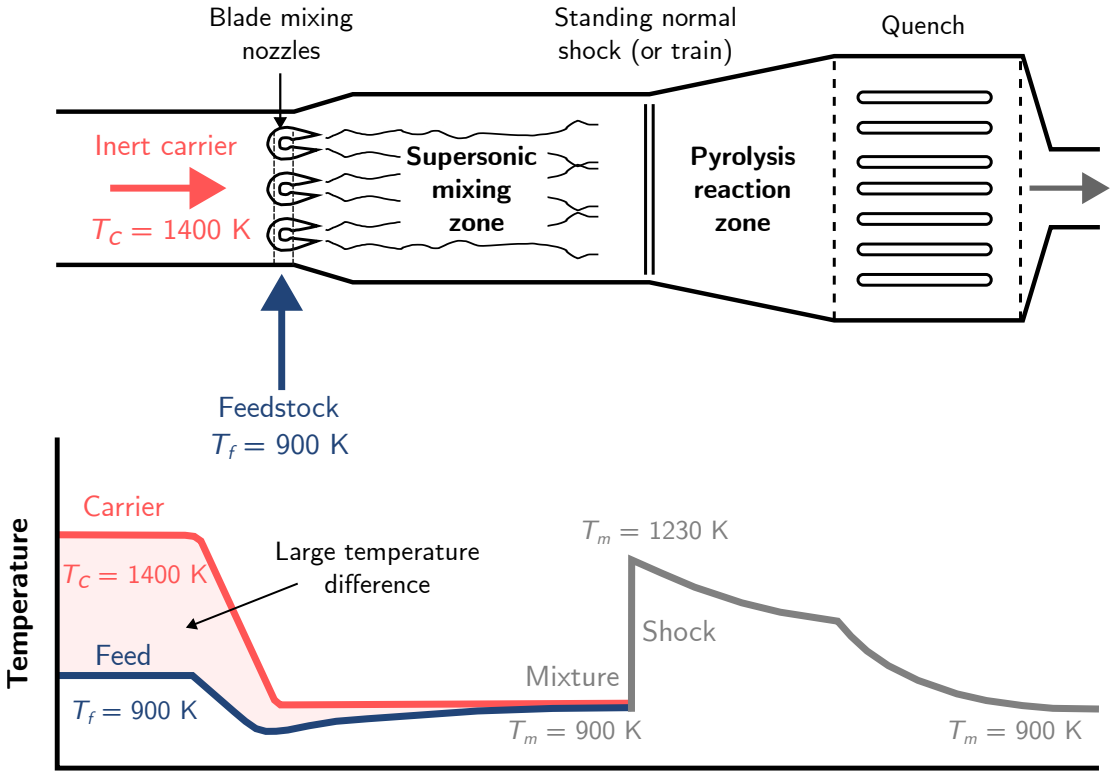


Figure 2.9: A schematic of the shockwave reactor used for the pyrolysis of hydrocarbons (reworked based on [84, 85])

However, the shockwave reactor failed to reach the market. It is not clear why; however, there are a few possible reasons. First, the device is still effectively a heat exchanger, a direct-contact heat exchanger in this case. Therefore, high temperatures are still needed to transfer energy to the inert carrier gas, which presents material challenges. Second, there are potential challenges with high energy consumption. To transfer sufficient energy between the inert steam carrier and the feedstock, a higher steam dilution ratio is required in comparison to conventional designs. This steam is

inert and does not participate in the chemistry; therefore, the energy consumption per tonne of high-value product produced increases.

Third, fundamentally, there is limited control over the reaction temperature profile since energy is converted almost instantaneously through a single shock. Therefore, the temperature decays (see Fig. 2.9) due to endothermic heat absorption, which may not be optimal for many feedstocks [23, 84]. Finally, concerning our objective of decarbonisation, this concept does not provide a solution. This is because the carrier gas must still be preheated in a furnace before entering the SWR.

Now, with a new drive to modernise furnace technology due to the demands of decarbonisation, the notion of leveraging shockwaves to increase reaction performance has resurfaced in a different form in the turbo-reactor design. This concept addresses many of the failures of the SWR.

Advantages of the RDR Concept

A combination of ideas from supersonic compressors, impulse turbines, regenerative machines and SWRs have inspired the design of the turbo-reactor, culminating in four vital benefits compared to the technologies summarised in Sec. 2.1.

1. ***Higher energy conversion efficiency:*** High energy conversion efficiency from a renewable energy source to high-temperature heat. There are three sources of energy loss: an electrical transmission loss of 6%, electrical motor loss of 3 – 7% at rated load [86], bearing loss of 1%, and heat loss through the insulated casing and along the shaft of 1 – 5%. If the renewable energy source is located near the industrial plant, transmission losses can be eliminated, and the overall energy conversion efficiency could be as high as 91%.

2. ***Higher reaction efficiency:*** Because the metal temperature level is lower than that of conventional furnace tubes, higher fluid temperatures are possible within the metallurgical limits of the blade materials. Higher fluid temperatures, coupled with shorter reaction time-scales, facilitated by the supersonic throughflow speeds possible in a turbomachine, enable 10 – 20% higher reaction efficiency than a

conventional furnace [14]. If the RDR were deployed globally, this could lead to 80 – 170 TW h lower energy consumption and between $15 billion and $30 billion lower feedstock expenditures for the olefin production sector per year.

3. ***Lower coking:*** Due to the high wall temperatures in conventional hydrocarbon cracking, secondary reactions decompose olefins into solid carbon (*i.e.,* coke) that forms directly on the walls. The wall temperature in the turbo-reactor is 100 – 200 °C lower than in conventional furnaces, and since coking is a strong function of temperature, the deposition rate is dramatically reduced. As a result, plant availability improves and carburisation reduces, which can substantially improve service life.

4. ***Scalability:*** Ease of scalability due to a large body of inherited know-how from high-temperature and high-capacity industrial turbocompressors, gas turbines and steam turbines. The turbo-reactor reduces the furnace footprint by 50 to 100 times, mitigating some challenges to scaling the concept and significantly reducing capital, operating and maintenance costs. Furthermore, cooling can be used to mitigate material challenges, oxidation and coking.

None of the other pathways discussed in the previous section (Sec. 2.1) simultaneously offer all the advantageous features that the RDR provides. These include improvements in capital cost, plant footprint, reaction efficiency and plant availability, which should help counteract high electricity prices relative to gas prices. Therefore, not only does the turbo-reactor achieve decarbonisation, but it is also a financially attractive solution thanks to innovations in the quality of the process.

The next section will introduce the fundamental principles and flow physics of the elemental stage design, and demonstrate using high-fidelity simulations the feasibility of using a novel turbomachine to induce endothermic cracking reactions. This section concentrates on the robustness of the *regenerative* multi-stage architecture; however, the multi-stage machine has also been configured as an axial (see Chapters 3 & 4), radial or mixed-flow type depending on the customer, the process and/or throughflow capacity requirements.

Pages 32–48:

A New Robust Regenerative Turbo-reactor Concept for Clean Hydrocarbon Cracking

(*J. Glob. Power Propuls. Soc.* 2022)

GPPS | JOURNAL

JOURNAL OF THE GLOBAL POWER AND PROPULSION SOCIETY
journal.gpps.global

A new robust regenerative turbo-reactor concept for clean hydrocarbon cracking

Original article

Article history:
Submission date: 19 May 2022
Acceptance date: 30 May 2022
Publication date: 28 July 2022

This is the updated version of a paper originally presented at the Global Power and Propulsion Virtual Technical Conference, GPPS Xi'an21, April 11–13, 2022.

Check for updates

***Correspondence:**
DR: dylan.rubini@eng.ox.ac.uk
NK: nikolas.karefyllidis@eng.ox.ac.uk

Peer review:
Single blind

Copyright:
© 2022 Rubini et al. This is an open access article distributed under the Creative Commons Attribution Non Commercial No Derivatives License (CC BY-NC-ND 4.0). Unrestricted use, distribution, and reproduction of the original work are permitted for noncommercial purposes only, provided it is properly cited and its authors credited. No derivative of this work may be distributed.

Keywords:
hydrocarbon cracking; decarbonisation; high-temperature processes; turbomachinery; aerodynamics; LES

Citation:
Rubini D., Karefyllidis N., Xu L., Rosic B., and Johannesdahl H. (2022). A new robust regenerative turbo-reactor concept for clean hydrocarbon cracking. Journal of the Global Power and Propulsion Society. 6: 135–150.
https://doi.org/10.33737/jgpps/150550

OPEN ACCESS

Dylan Rubini[1,*], Nikolas Karefyllidis[1,*], Liping Xu[2], Budimir Rosic[1], Harri Johannesdahl[3]

[1]*Oxford Thermofluids Institute, University of Oxford, Oxford OX1 3PJ, UK*
[2]*Whittle Laboratory, University of Cambridge, Cambridge CB3 0DY, UK*
[3]*Coolbrook Oy, Pieni Roobertinkatu 9, 00130 Helsinki, Finland*

Abstract

Energy consumption in the chemical and petrochemical industries is dominated by the highly energy-intensive, high-temperature hydrocarbon cracking process. This paper offers an attractive solution to potentially decarbonise the cracking industry by direct renewable electrification whilst also improving the efficiency of energy transfer and enhancing the quality of the overall cracking process. Controlled endothermic cracking reactions are generated by direct mechanical energy transfer using a new electric-motor-driven turbo-reactor to replace the fossil-fuelled radiant unit in the conventional cracking plant. This study provides an introduction to the novel concept, focusing on how the unique three-blade-row repeating stage design is exploited in order to drive the efficient energy transfer and transformation processes used to overcome the fundamental limitations of the conventional process. More specifically, this includes increasing the primary product yield, mitigating coke deposition and improving operability. A combination of URANS and high-fidelity LES are presented in order to probe the aerothermal flow physics and interactions in the rotor and diffuser. This study clearly demonstrates how the working principals of the elemental stage can be realised uniformly across the full regenerative turbo-reactor by exploiting the robust and naturally self-adjusting concept. This work highlights the high-level of controllability of the turbomachine, which allows the cracking environment and species selectivities to be tailored during operation, for example, by conducting minor alterations to the rotational speed.

Introduction

Over 10% of the global final energy demand is attributed to *high-temperature* ($T > 400°C$) industrial fossil-fuel-based heat requirements alone (IRENA, 2016; de Pee et al., 2018). In contrast with low-temperature heat, which can be decarbonised with more mature technologies such as heat pumps, economically and technically feasible technologies aimed at decarbonising high-temperature industrial heat are limited. One vital industry requiring urgent innovation is the highly energy-intensive steam cracking industry, where ~110 million tonnes of natural gas (equivalent to ~4% of global natural gas consumption from all sectors) are consumed annually (Mynko et al., 2022). As a result, a staggering ~300 million tonnes of CO_2 are emitted per year (Keller et al., 2020), accounting for 30% of global emissions from the chemical and petrochemical industries.

This paper presents a new turbo-reactor as an economically and technically feasible solution which could potentially decarbonise the hydrocarbon cracking industry. Figure 1 shows the novel concept which

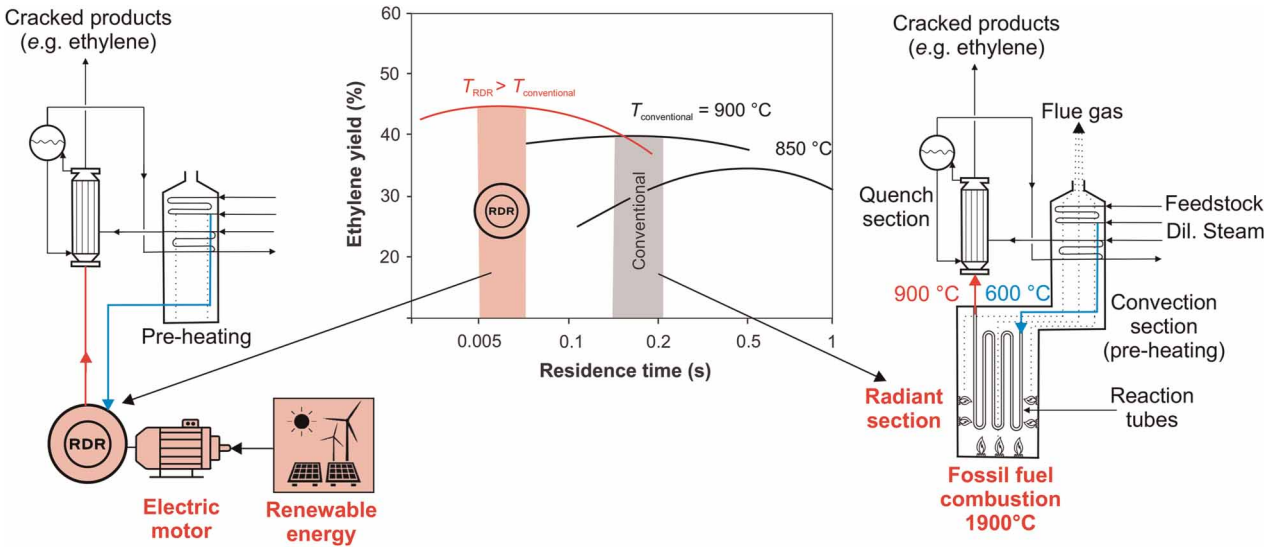


Figure 1. A schematic of an ethylene production plant including the turbo-reactor (left) and a conventional steam cracking furnace (right). The middle schematic shows how shorter residence times higher temperatures can increase the yield (adapted from Chauvel and Lefebvre (1989)).

eliminates the fundamental limitations of the conventional process by replacing the inefficient, ineffective and indirect surface heat exchange energy transfer mechanism in the radiant section with efficient and *direct* mechanical energy transfer using an electric-motor-driven turbomachine.

Based on the levelised cost of electricity, grid-parity is being achieved for utility-scale solar PV as well as onshore wind (IRENA, 2014). Coupled with wide-spread carbon tax deployment in the near future, this could mean that powering the cracking process using *cheap* renewable electricity will become a significantly more attractive solution than natural gas combustion within a furnace. Complete renewable electrification of this unit would decrease CO_2 emissions by at least 70% for a steam cracking plant. Most importantly, Figure 2

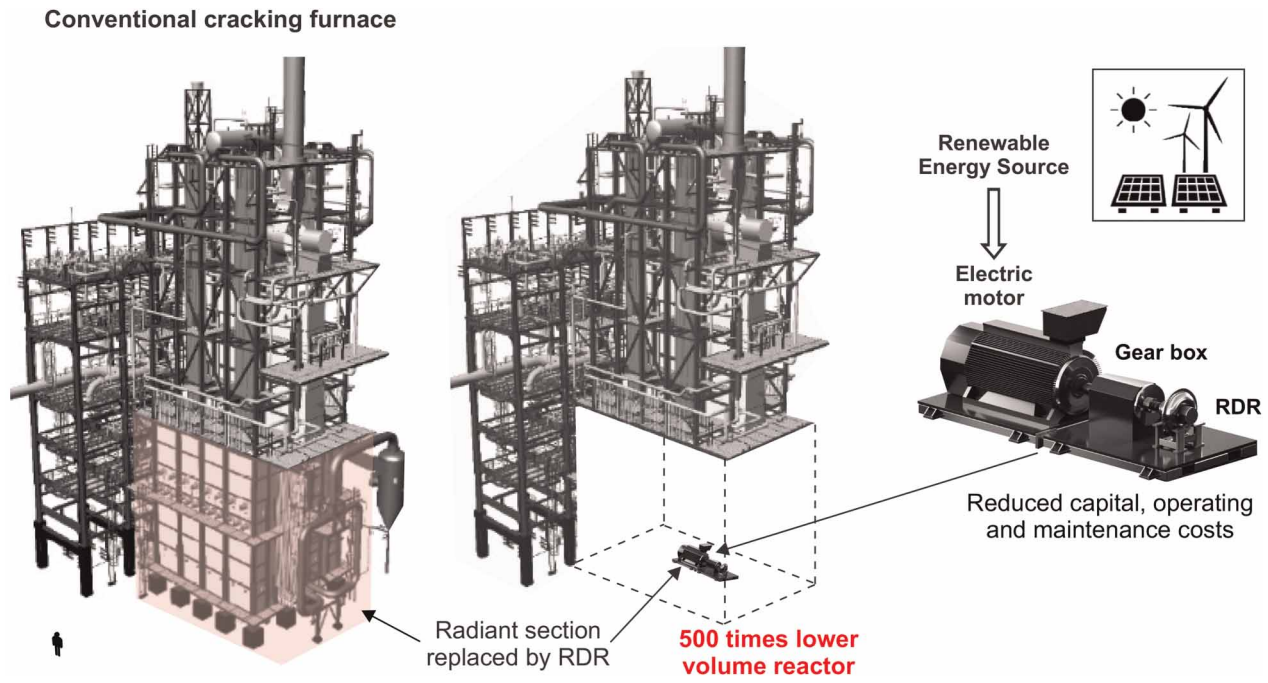


Figure 2. A size comparison between the radiant section (adapted from Linde Engineering (2020)) and a scaled representation of the RDR within a steam cracking plant.

demonstrates that the new turbo-reactor provides the necessary enthalpy rise for the endothermic cracking reactions within a low-cost, single-rotor, and compact machine with a volume approximately 500 times smaller than a conventional furnace. This on its own is a potential step-change improvement for the cracking industry.

The new turbo-reactor, the Roto-Dynamic Reactor (RDR), not only enables a more efficient low-carbon energy transfer and transformation process, but also has the potential to substantially improve the quality and controllability of the overall cracking process. This includes a projected increase in the desired primary product yield (see Figure 1) as well as a substantial reduction in coke formation and deposition rates. One important mechanism to accomplish this is a controlled residence time reduction (see Figure 1 (middle)), which is achieved using energy transfer and transformation processes of a shorter time-scale and higher reaction temperature. This is enabled by exploiting an *ultra*-high loading rotor and a shockwave system.

Alternative zero-carbon technologies which could replace the conventional furnace include zero-carbon fuel-switching (e.g., hydrogen (Haydock and Napp, 2013)) or the use of electric furnaces (Amghizar et al., 2020). However, these technologies do not overcome the fundamental limitations of the (exergetically) inefficient and ineffective surface heat exchange energy transfer mechanism, nor do they enhance the quality of the process.

This numerical study illustrates how the working principals of the single elemental stage (Rubini et al., 2021) can be realised in the context of a regenerative multi-stage environment. The robust concept ensures the required conditions for high-quality cracking are maintained for all stages of the turbo-reactor, whilst using a simple, uniform, compact, and low-cost design. In contrast with conventional reaction tubes, the turbomachine is highly controllable, allowing the level of specific energy imparted, temperature profile, pressure profile and mixing intensity to be tuned to provide the necessary cracking severity during operation.

Conventional hydrocarbon cracking process

The heart of the cracking process is the radiant section (see Figure 1 (right)) which provides the large enthalpy of reaction to thermally decompose (crack) a wide variety of long-chain hydrocarbon feedstocks (e.g., naphtha or bio-feeds) into a broad range of short-chain high-value light olefins (e.g., ethylene). These final products are essential building blocks for a wide range of industrial and domestic products. The radiant furnace is the most important unit of the plant and is the section that the new electric-motor-driven turbo-reactor has been developed to replace (as shown in Figure 1 (left)).

Before the feedstock enters the radiant section, it is mixed with steam in order to reduce the olefin partial pressure and therefore minimise the level of yield-degrading secondary reactions. Following the blue line through the convention section shown in Figure 1 (right), the gas mixture is preheated by heat exchange with the flue gas (the dotted black line originating from the radiant section). In the radiant section, the radiative heat flux from the combustion of fossil fuels (e.g., natural gas) drive the endothermic cracking reactions as well as raising the temperature of the gas mixture from ~600°C to ~900°C along the reaction tubes. Downstream of the radiant section, the gas mixture is rapidly cooled in the quenching unit in order to freeze the gas composition. The cracked product mixture (primary ethylene and propylene alongside secondary products) is compressed, chemically treated, and separated into the desired final products.

Addressing the fundamental limitations of the process using a new turbo-reactor

The primary product yield and the level of coke deposition are two key metrics defining the quality of the process. High fluid temperatures (see Figure 1), low metal temperatures, short residence times, as well as a constant olefin partial pressure are the key drivers which can increase the yield and curtail coke deposition. These parameters have limited scope for improvement in conventional designs. The proceeding section highlights how the new turbo-reactor removes the fundamental limitations of the conventional furnace to provide an efficient energy transfer mechanism whilst at the same time substantially improving the effectiveness and quality of the cracking process.

The Roto-Dynamic Reactor (RDR) concept and detailed design have been developed over many years of collaboration with Coolbrook Oy (2020) and are based on the original patent by Bushuev (2016).

Efficient energy transfer as work input for the chemical reactions

Conventional process limitation. The prevailing radiative heat transfer mechanism in the radiant section (where high heat fluxes are required) necessitates very large temperature differences to drive energy transfer. Therefore,

the exergy destruction here dominates and contributes to a deleteriously low rational exergy efficiency of 43% for the furnace (Yuan et al., 2019).

Turbo-reactor solution. Mechanical energy (i.e., pure exergy) is imparted directly by the aerodynamic action of a rotating blade row. This eliminates the demand for large temperature differences within the system to drive the heat fluxes in the energy transfer process. As a result, the energy is directly imparted into the reactants without significant heat loss, especially with the application of effective modern thermal protection to minimise irreversible heat losses. It will be demonstrated that the required enthalpy can be provided using only a single-rotor blade row within a ~500 times smaller volume reactor than a conventional furnace. It is also noted that the operating temperature ($T_{max} \approx 900 - 1050\,°C$) and pressure ($p \approx 1.5 - 3.0$ bar) are less severe than those experienced in conventional gas turbines; therefore, available and mature material technology can be used.

Increasing the yield though high gas temperatures and short residence times

Conventional process limitation. Indirect surface heat exchange in the radiant section requires the walls to be at a significantly higher temperature than the gas mixture within the reaction tubes, often by 150°C (Van Goethem, 2010). Therefore, the working fluid temperature is fixed by the metallurgical constraints of the tubes as well as by the acute coke deposition rates associated with the elevated metal temperature. Figure 1 illustrates that this fixes the optimum residence time required to maximise the yield at a given temperature level.

Even supposing higher gas temperatures were possible, the limitations of the conventional design make it challenging to reduce the residence time to the optimum value required to maximise the yield at the higher temperature profile (see Figure 1). This is because the fluid velocity is fixed by the pipe geometry and pressure losses, and the gas path length is fixed by the necessary heat transfer surface area.

Turbo-reactor solution. In contrast with a conventional furnace, the average metal temperature is lower than that of the working fluid. Therefore, higher gas temperatures can be attained without the walls reaching the critical yielding value, whilst at the same time alleviating coke deposition. In conjunction with this, lower residence times are now feasible, firstly, because the requirement for a large heat transfer surface area is removed (thus allowing a shorter gas path), and secondly because the fluid can reach significantly higher throughflow speeds than within a pipe. Figure 1 indicates that these factors lead to the necessary conditions to increase the yield by reducing the residence time to match the increased temperature.

Potential to reduce coke formation and deposition

Conventional process limitation. Coke formation and deposition is the primary challenge impacting plant availability, operating and maintenance costs, and component lifetime. Generally, coke deposition is magnified by high wall temperatures and low surface viscous shear stresses, as well as long residence times and elevated olefin partial pressure.

Coke is continually deposited on the inner wall surface and accumulates over time. This leads to an effective thermal resistance up to 80% higher than for a bare tube (Zhang et al., 2022). To maintain the same cracking severity and product yield, the resulting increase in thermal resistance is compensated for by increasing the fuel flow rate to the burners, therefore raising the tube metal temperature and further exacerbating coking. As a result, coke deposition provides a stringent restriction on the allowable gas temperatures within the tubes which thus limits the yield. Coupled with this, coking reduces olefin selectivity by increasing the pressure drop, as well as reducing component lifetime by increasing thermal stresses. After several weeks, coke can occupy up to ~50% of the cross-sectional area of the reaction tube (Vandewalle et al., 2017). Therefore once a critical tube metal temperature or pressure drop is reached, the furnace is taken completely offline for a 48 h decoking period which occurs as frequently as every 6 weeks (Schietekat et al., 2014) and adversely affects the profitability of the plant.

Turbo-reactor solution. The new turbo-reactor provides the necessary conditions to potentially reduce coke formation and deposition. This is achieved through a significantly lower wall metal temperature and by eliminating the outward wall temperature gradient (from metal to fluid), in addition to shorter residence times.

Controllability

Conventional process limitation. Due to the nature of the fixed pipe geometry in conventional reaction tubes, the only available control parameters are the initial temperature and pressure of the feedstock and the amount of the fuel supplied to the burners. Accurate control over the temperature, pressure, and residence time distributions along the tubes is challenging.

Turbo-reactor solution. Through careful aerothermal design and operational control (e.g., varying the rotational speed and mass flow rate), turbomachines have a high level of control over the temperature and pressure distributions, specific enthalpy increase, as well as the chemical reaction rate, residence time, turbulent mixing of species, and the steam feed injection staging. This enables controllable tailoring of the olefin selectivity and product yield.

The turbo-reactor concept

Three-blade-row elemental stage and regenerative heating concept

Based on the design requirements necessary to provide the optimum conditions for cracking, the repeating elemental stage consists of a novel three-blade-row design including stator, rotor and diffuser blade rows (see Figure 3a).

The stator blades (inlet guide vanes) are designed to accelerate, pre-condition and pre-swirl the flow. The proceeding sections highlight how a controllable optimisation of the residence time (and hence yield) is enabled using a transonic ultra-high loading rotor and a complex shock system in the diffuser blade row. This facilitates ultra-short-distance energy transfer and transformation processes. The transonic diffuser is multi-functional, but its primary function is to turn and slow down the high-speed flow, as well as transform kinetic into internal energy for the reaction whilst also maintaining a constant static pressure level (optimum for cracking) downstream.

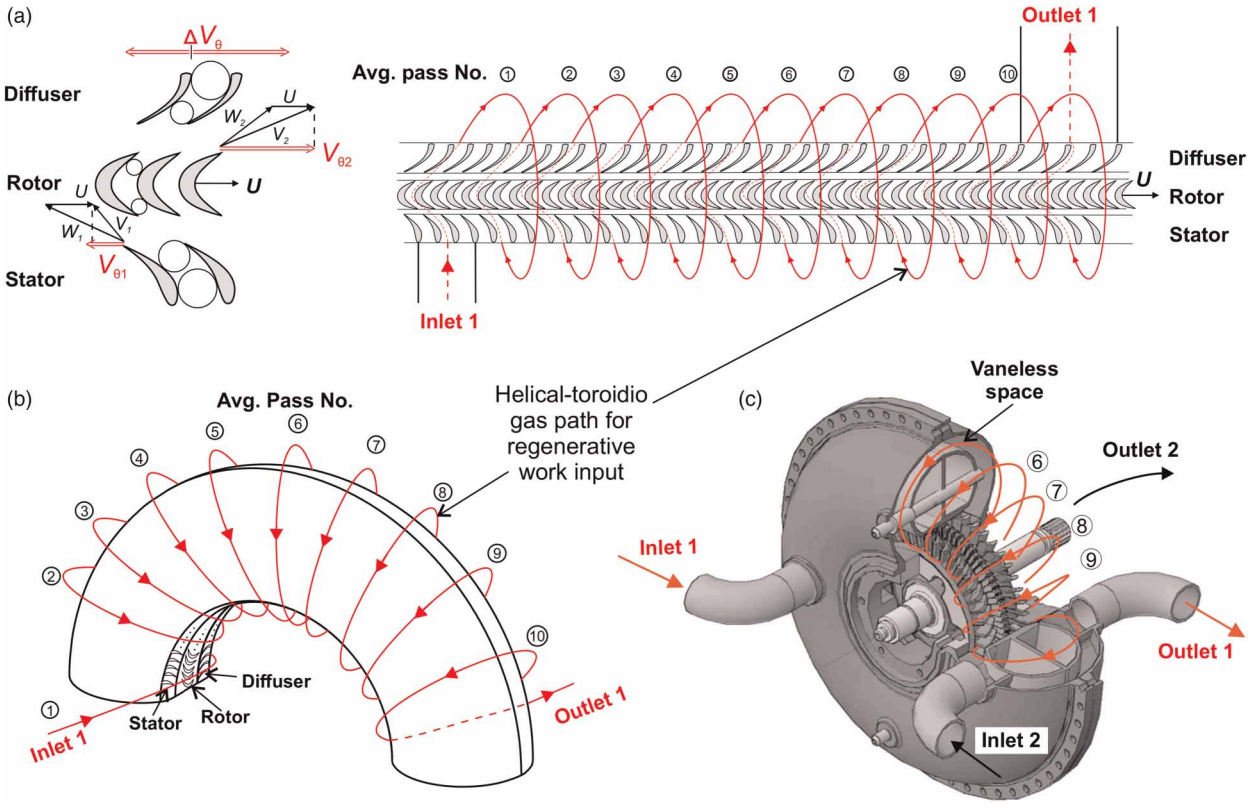


Figure 3. (a) A linearised schematic showing the hemi-annulus and the blade velocity triangles (b) The Hemi-annulus schematic (c) The full reactor model.

The work input across a single elemental stage does not provide a sufficient total enthalpy rise to complete the cracking reactions, therefore, a multi-stage configuration is necessary. In this version of the turbo-reactor, this is achieved regeneratively using only a single rotor blade-row. Figure 3 shows that a series of regenerative passes (indicated in red) are used to compound the total enthalpy rise across the machine for the chemical reaction. Following the red line, a single regenerative pass constitutes an average particle traversing the stator, rotor and diffuser blade-rows, before following a helical-toroidio gas path in the vaneless mixing/reacting space. The gas path indicated in red in Figure 3 illustrates that the diffuser trailing edge (TE) allows sufficient residual swirl to progress the flow through the vaneless space to re-enter the next stator passages shifted in the circumferential direction. The regenerative multi-pass design has a similar effect to multiple turbomachinery stages in series. The vaneless space is formed between the inner and outer casing, as shown in Figure 3c. It is noted that the helical-toroidio regenerative flow path is similar to that in a regenerative compressor (Sixsmith and Altmann, 1977).

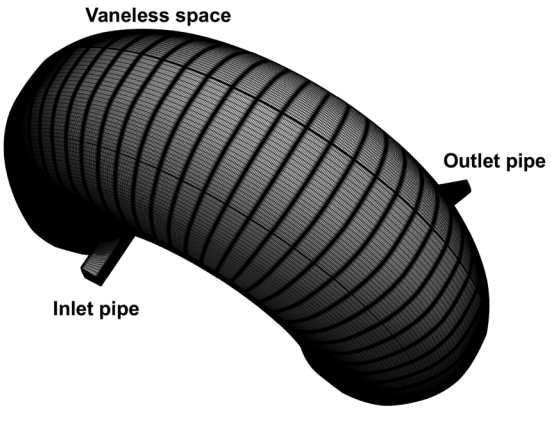
The regenerative process is repeated for up to 10 designed passes of repeated turns before exiting through the outlet pipe. Figure 3c shows a second independent but identical hemi-annulus which is mirrored in the lower half of the full reactor, resulting in two inlet and two outlet ports.

The unique advantage of the regenerative turbomachine is that it uses only a single rotor with a short shaft length, resulting in a compact low-cost reactor ideal for small-scale production. Various complementary architectures are being investigated for large-scale production capacity. Nonetheless, for increased capacity a multiplicity of these compact regenerative machines can be run in parallel.

Numerical methodology

The feasibility of the concept has been proved with a numerical model simulated within the in-house computational fluid dynamics code TBLOCK developed by John Denton (Denton, 1982, 1986, 1992). The solver has been applied to a wide range of flow physics within as broad range of turbines and compressors test cases, as well as being extensively validated against experimental data. The experimental validation of the solver most relevant to the study presented here is the work of Rosic et al. (2005). This work confirmed the capabilities of TBLOCK against experiment data for a multi-stage turbine with a Reynolds number similar to that in the new turbo-reactor (see the table in Figure 4). In addition to URANS calculations, high-fidelity large eddy simulations (LES) have been employed to more accurately capture the mixing-out of large-scale coherent turbulent flow structures and improve the resolution of shockwave boundary layer interaction.

Different subgrid scale (SGS) models have been implemented in the solver by Klostermier (Klostermeier, 2008). The SGS model used for the present computation is the Smagorinsky model coupled with the Van Driest damping function. The present work adopts an adiabatic wall as the thermal boundary condition for the energy equation. Additionally, a semi-perfect gas consisting of a single species with constant chemical composition is assumed. The gas composition is very close to the average of the inlet gas mixture (naphtha and steam) and the cracked product mixture at the outlet. In the CFD model presented, there is no chemical reaction model yet as the primary focus of this preliminary design is on the aero-thermodynamics of the turbo-reactor.



Operating conditions	
U (ms^{-1})	280
Inlet T_0 (K)	823
Inlet p_0 (bar)	2.0
Exit p (bar)	2.0
Re	50,000 – 200,000

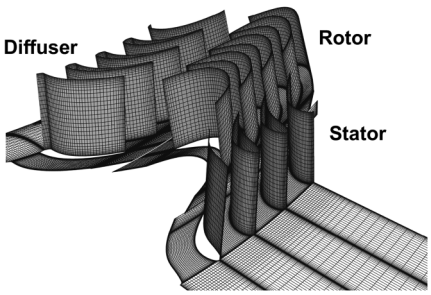


Figure 4. Computational mesh and operating conditions for the URANS calculations.

Table 1. Mesh sensitivity study.

	Coarse	Medium	Fine
No. grid nodes (bladed path)	48 × 32 × 32	81 × 51 × 51	91 × 71 × 71
Stage loading ψ	4.05	4.10	4.07
Flow coefficient ϕ	0.6009	0.6054	0.6033

Computational domain

The URANS domain shown in Figure 4 consists of 527 blocks with 48 × 32 × 32 nodes in the axial, pitchwise and spanwise directions within blade passages (totalling to 16×10^6 grid points). Turbulence is modelled using wall functions and an algebraic mixing length model.

The justification for the mesh density which has been used for the URANS calculations is two-fold. Firstly, the required mesh density is a function of the turbulence closure model, which in this case is a mixing length model. Therefore, the sensitivity to mesh density is minimised. Secondly, the mesh density is dictated by the physical phenomena of interest. In this study, only zeroth-order parameters of the elemental stage are of primary concern. Namely, the present work aims to confirm that the stage achieves the required work input (enthalpy gain) at the flow coefficient that is necessary to maintain the required throughflow capacity. These bulk flow parameters are relativity insensitive to the mesh density. Since flow coefficient and work coefficient are the primary parameters of interest, they are used to define the objective function used for a mesh sensitivity analysis. For investigations into detailed flow physics and flow structures, the mesh-dependency is greater. However in such cases, LES is used in this study.

For consistency, during the mesh sensitivity study the wall expansion ratio is kept constant and the number of grid points in each direction of the computational space are scaled equally. Table 1 illustrates that the maximum variation in stage loading and flow coefficient is about 0.5% across the different mesh resolutions. Relatively low deviations are expected since the flow in the blade passages is predominantly pressure driven. Similar deviations are observed for other flow parameters of interest, such as pressure, internal energy and dissipation rate along the gas path.

The regenerative nature of the process allows for a truncated domain (consisting of a single pass or streamtube to be modelled when conducting the LES (see Figure 5a). The exit static pressure boundary condition ($p = 2.1$ bar) is placed at the end of the vaneless space based on a precursor URANS simulation of the full hemi-annulus. A periodic boundary condition is applied to the lateral surfaces of the domain. The LES domain consists of 48 blocks with 241 × 131 × 131 nodes in the axial, pitchwise and spanwise directions within the blade passages (totalling 150×10^6 grid nodes across the domain), resulting in a time-averaged y^+ between 0.1 – 1.0. The energy cascade shown in Figure 5b suggests that the LES is capable of resolving the time scales of the energy carrying

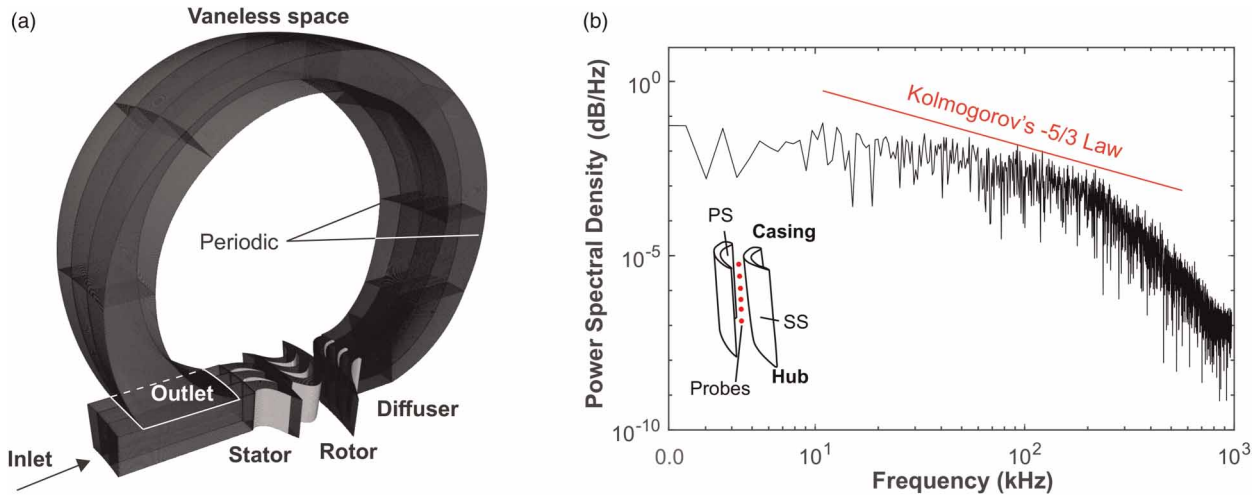


Figure 5. (a) LES computational domain (b) Turbulent kinetic energy cascade averaged over six diffuser stations.

eddies in the inertial sub-range of the Kolmogorov energy spectrum, strengthening the validity of the numerical results.

Results and discussion

Short-distance work input: ultra-high loading rotor

The novel design requirements of the new turbo-reactor require high work input and energy transformation rates in order to controllably optimise the residence time and thus maximise the yield. These conditions are provided by the energy transformation train illustrated in Figure 6. One key component of the energy train is the ultra-high loading transonic rotor, which is used to impart the required energy into the working fluid over an extremely short distance and time scale. For a given mass flow, the primary objective of the turbo-reactor is to maximise the amount of energy imparted, which can then be dissipated into internal energy to drive the reaction. Therefore, stage loading is used as a metric to quantify aerothermal performance.

Based on the specific design objectives and operating conditions, a rotor stage-loading coefficient ($\psi \approx 4.0$) an order of magnitude higher than a typical axial compressor is possible. Therefore, a very high level of specific energy (≈290 kJ kg^{-1}) can be imparted into the feedstock per stage. Figure 6 illustrates that this is achieved using very high turning symmetric blading, supersonic (relative) ramming at the rotor inlet, and supersonic (absolute) Mach numbers at the rotor outlet. The majority of the specific energy imparted by the rotor is in the form of kinetic energy which must then be converted, in the downstream diffuser and vaneless space, into internal energy for the reaction.

Short-distance energy transformation: diffuser shock system

In contrast with the more conventional stator and rotor blade rows, the diffuser is multi-functional with three primary objectives which are highlighted in Figure 6.

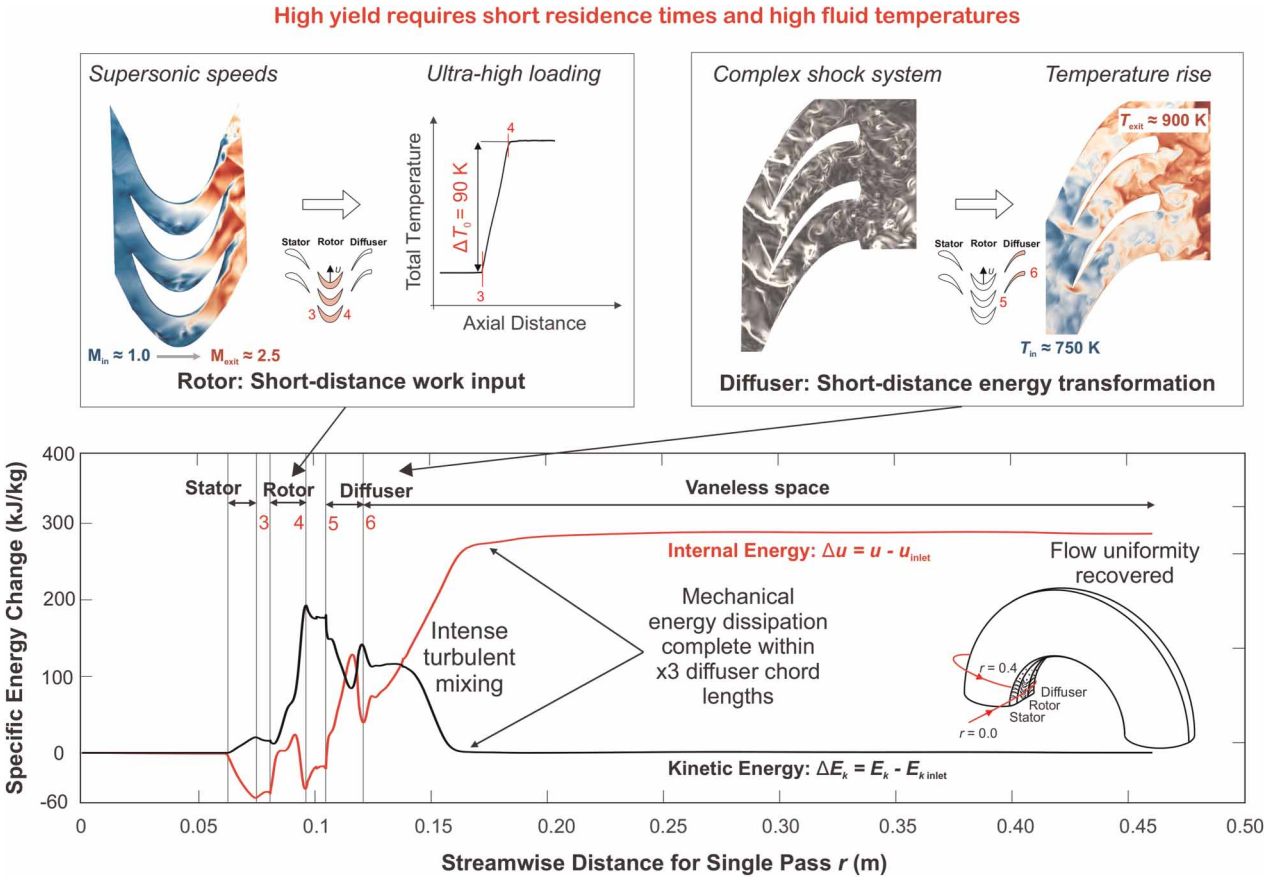


Figure 6. Energy transformation train along a single gas path: time- and mass-averaged specific internal u and specific kinetic energy E_k changes referenced as a change relative to the inlet values (LES).

Firstly, the diffuser transforms the bulk kinetic energy of the flow provided by the rotor into internal energy for the reaction. This is achieved using a shock train system to slow the high-speed flow down within the passage. Secondly, the diffuser is designed to promote large-scale flow separation and mixing in passage in order to prevent pressure accumulation whilst dissipating the kinetic energy, primarily through intense shockwave boundary layer interaction (SWBLI) and shock-induced flow separation (see Figure 6). This is necessary in order to maintain a constant static pressure level which maximises the olefin production rate (by reducing side reactions) and promoting homogeneity in the temperature and species field. Finally, as a consequence of highly impulsive rotor blading, the diffuser must turn the flow through a large angle (almost back into the axial direction) in order to allow the flow to progress regeneratively downstream of the diffuser TE. The high-turning blade profile further promotes shockwaves in the passage as well as augmenting the mixing levels through increased flow separation.

Figure 6 (bottom) shows that the energy transformation processes are complete within approximately three diffuser chord lengths downstream into the vaneless space, before the flow conditioning process commences in order to recover flow uniformity.

Controlled turbulent mixing and flow conditioning processes

The vaneless mixing/reacting space shown in Figure 7a has two primary functions and can be split into two distinct zones: the primary mixing/reaction zone and the flow conditioning zone. As a result of a thick diffuser TE, high-energy turbulent mixing, driven by large coherent vortical structures (e.g., vortex shedding), propagates into the vaneless space. The controlled mixing process generates the ideal conditions for cracking by maintaining a constant static pressure level, homogenising the temperature field, converting kinetic into internal energy, and possibly accelerating the chemical kinetics.

Within three diffuser chord lengths, due to the low Reynolds number regime (50,000–200,000) throughout the vaneless space, flow uniformity is rapidly recovered (see Figure 7) and the pressure is relieved back to the inlet level. Detailed studies (Rubini et al., 2021) have shown that the turbulence intensity decays rapidly downstream of the high-energy mixing zone. In addition, the momentum thickness Reynolds number and shape factor, two parameters which characterise the state of the boundary layer, both recover back to similar values to those of the clean flow within the inlet pipe.

Figure 7b shows the time-averaged footprint of large length-scale vortical structures which are stretched as the flow regains uniformity. Any vortical structures or turbulent fluctuations that persist before the next pass will be stretched and dissipated by rapid acceleration in the stator. It can also be seen that in the majority of the vaneless space, the level of lateral mixing between passes is low, which is important for the quality and effectiveness of the regenerative design. Before the next rotor stage, the flow is in a well conditioned state which maximises the work input potential of the following stage. Figures 6 and 7 both indicate that the mixing and energy conversion processes are complete before 50% of the length of the vaneless space, suggesting that tuning the vaneless space length (and hence residence time and reactor volume) is possible without influencing the flow conditioning for

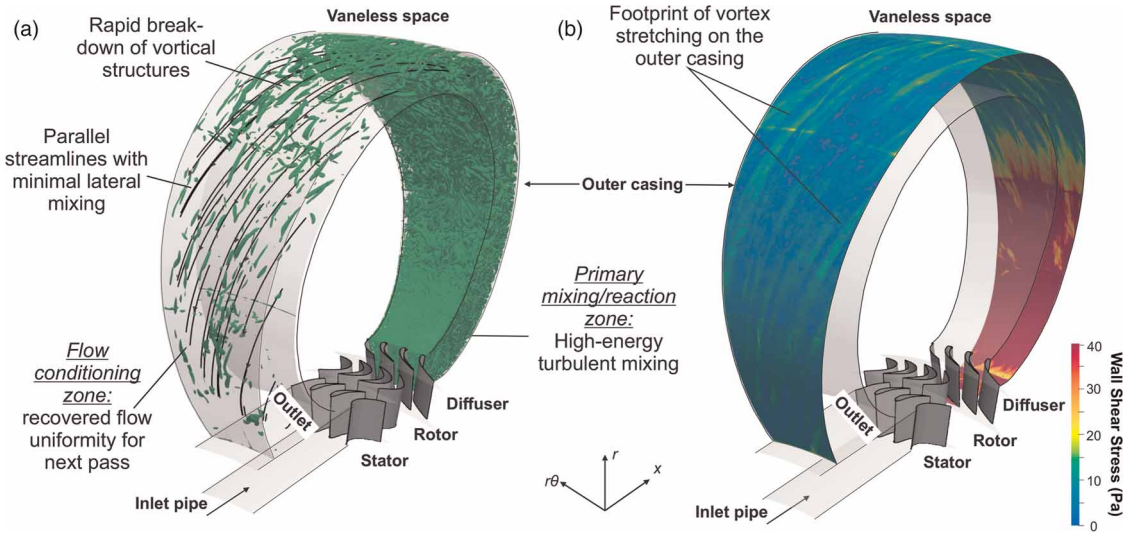


Figure 7. (a) Instantaneous (LES) Q-criterion iso-surfaces (b) Time-averaged (LES) wall shear stress contour on the vaneless space outer casing.

the next pass. This is further supported from a residence time control perspective since 80% of the total residence time is spent in the vaneless space.

Multi-pass regenerative energy input concept

Each successive regenerative pass repeats the processes described above for the elemental stage, resulting in a multi-pass regenerative heating concept as shown in Figure 8a. The following sections illustrate how the working principals of the elemental stage, described above, can be realised uniformly and effectively for all regenerative passes.

Figure 8a illustrates that there is sufficient tangential velocity at the diffuser TE to progress the flow regeneratively to the subsequent stator passages. For all regenerative passes, the streamtraces indicate that the flow is well-conditioned and that uniformity has been recovered downstream of the mixing zone. The open-volume vaneless space is robust enough to allow the streamlines to self-correct and form clearly defined passes without the need to partition the flow (see Figure 8).

Ideally, the streamtubes defining the passes must be independent in order to prevent lateral mass transfer of species between passes (which would degrade the performance of the reaction), preserve the required peak temperature for cracking, and maintain the design velocity triangles. The pattern of the streamtraces in Figure 8 indicate that the level of lateral mixing between adjacent regenerative passes is low. If necessary, this can be further controlled by partitioning each pass with thin guidewalls the vaneless space.

Figure 8b shows a uniform static enthalpy rise (relative to the inlet conditions) circumferentially across the reactor. For chemically reacting flow, depending on the dynamics of the reaction, the static enthalpy profile will be different to that shown in Figure 8b. Nevertheless, it has been clearly demonstrated that the necessary enthalpy rise for consumption by the endothermic cracking reactions can be provided within a compact, single-rotor turbomachine. This is achieved whilst maintaining a low wall temperature and mitigating coke deposition. The residence time for an average particle undergoing regenerative work input is 5–10 times lower than that for a conventional furnace. Coupled with the increased process gas temperatures, this could potentially increase the combined ethylene and propylene yield by 20% (Coolbrook Oy, private communication, 2019).

Robust aerothermal concept

The robustness of the aerothermal concept is enabled through the unique regenerative open-volume design. This robustness is characterised by two intrinsic features of the concept: a guaranteed pressure equilibrium, and the tolerance to density variations in spite of a low-cost, uniform casing and blade design.

Figure 9a illustrates that a constant static pressure level (p = 2.1 bar) is maintained along the turbo-reactor gas path as required by the reaction, indicating that the mixing and pressure equalisation is sufficient to prevent any

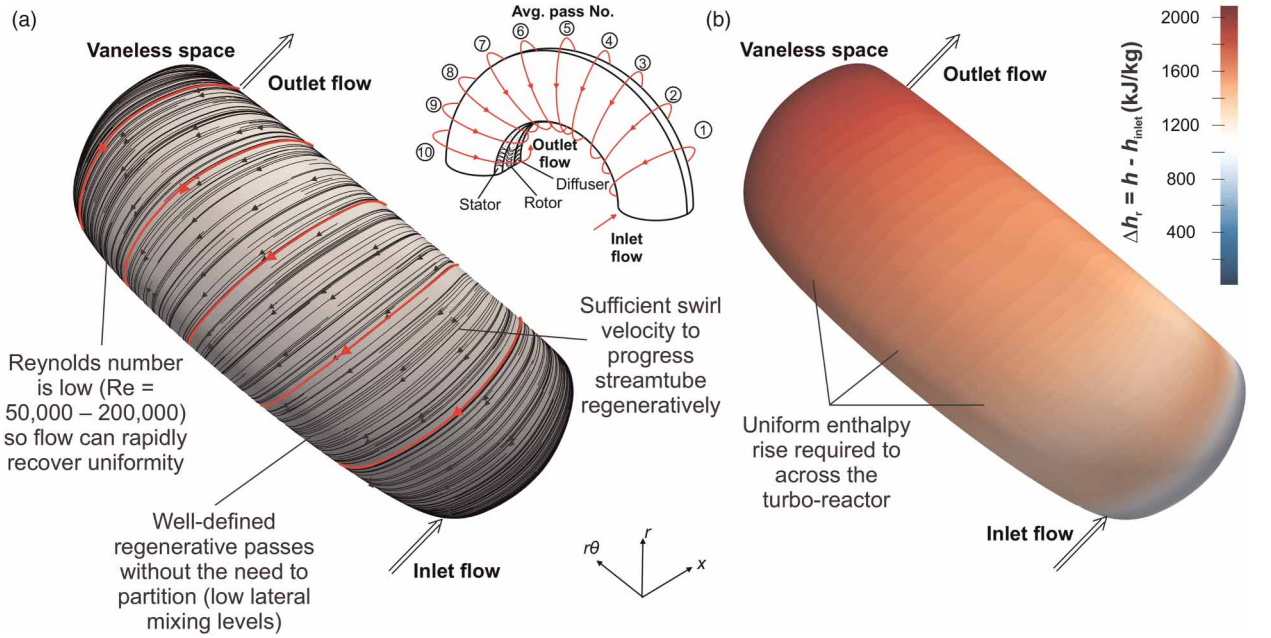


Figure 8. Time-averaged (URANS): (a) streamtraces (b) enthalpy change referenced relative to the inlet conditions.

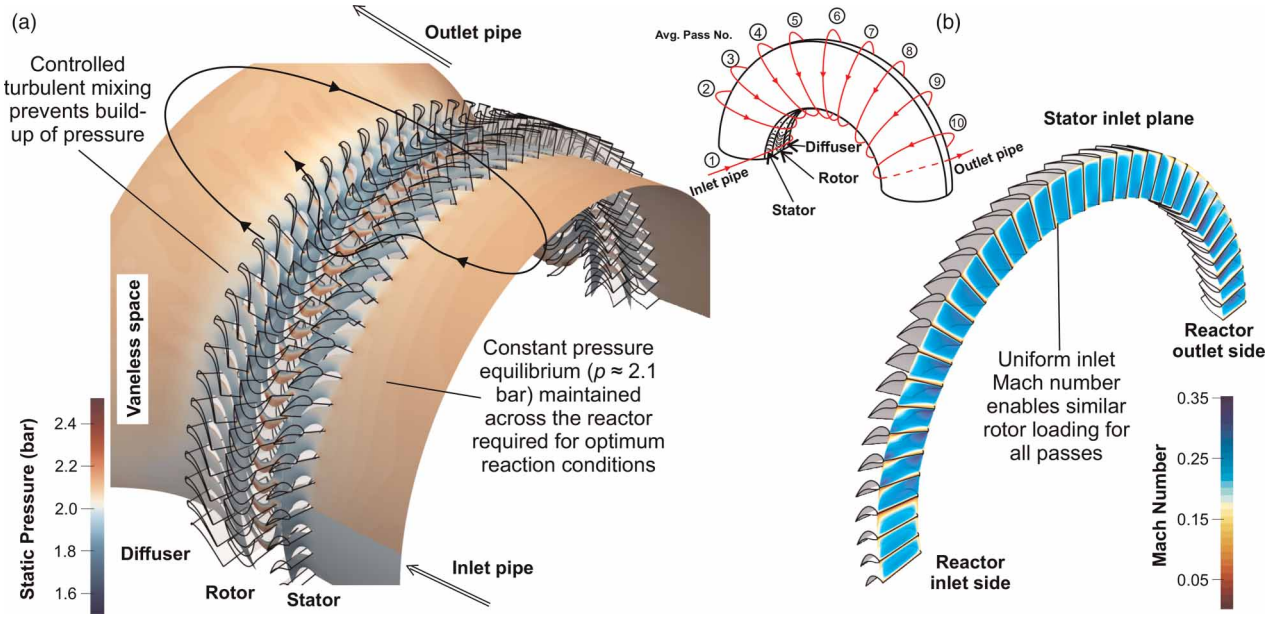


Figure 9. (a) Time-averaged (URANS) static pressure across the turbo-reactor (b) Time-averaged (URANS) Mach number at the inlet to the stator.

significant build-up of pressure. The time-scale associated with minor pressure spikes within the blade rows is is small, therefore the reaction is largely unaffected.

The advantage of using the unique open-volume vaneless-space is that it guarantees that a nearly constant pressure equilibrium is maintained across the reactor (see Figure 9a) as a result of the circumferential communication of pressure waves. The pressure level can be controlled with valves in the pipes. Within the open-volume vaneless space the flow naturally self-regulates in order to maintain well defined passes with straight and parallel streamlines (see Figure 8a). Therefore, this must be reflected by uniform pressure, which is a key characteristic defining the robustness of the concept.

From inlet to outlet, the increasing fluid temperature drives the density to decrease (by ~35%) without an associated change in pressure. The nature of the open-volume vaneless space allows the streamtubes to naturally expand in response to this density reduction. Since mass is conserved, this streamtube expansion aids in maintaining an almost uniform inlet Mach number and flow coefficient across the reactor despite the density changes (see Figures 9b and 10c). Accompanied with this, Figure 10b illustrates that a similar stator blade loading profile is sustained for all passes, once again indicating well conditioned flow for all passes and therefore highlighting the robustness of the concept. The stator is designed with a thick leading edge to provide tolerance to incidence variations. This ensures that the necessary tangential velocity is presented to the rotor despite minor distortions to the flowfield at the stator inlet (see Figure 10b).

The rotor stage loading coefficient is a direct function of flow coefficient and the stator exit flow angle (which can be inferred from stator blade loading profiles shown in Figure 10b). Therefore it has been demonstrated that despite density variations, the robust concept enables a simple, circumferentially uniform blade profile and inner/outer casing design to be applied whilst sustaining a high level of specific enthalpy input across all rotor stages (see Figure 10c).

For true chemically reacting flow, the gas composition changes (during cracking) from inlet to outlet, which augments the temperature-driven density reduction across the turbo-reactor. This is reflected by a decrease in the Reynolds number from 200,000 (inlet) to 50,000 (outlet). Despite this, simulations (not included here for brevity) have indicated that the dependence of the aerothermal flowfield on the Reynolds number is relatively weak within this range. Importantly, the change in the rotor blade loading (work input) and dissipation rate (mixing) for these two Reynolds numbers are within a few percent.

If necessary, it is possible to circumferentially vary the blade pitch of the stationary blade rows to compensate for the effects of density and gas property variations due to the reaction (not modelled) on other non-dimensional groups (e.g., Mach number). This would ensure that the blade loading (and hence exit flow angle) is uniform for all passes by maintaining a constant passage mass flow rate. Therefore, a high-level of specific energy input can be imparted across the turbo-reactor. However even with a uniform design, the robust open-volume space should partially tolerate these variations through a natural streamtube self-adjustment.

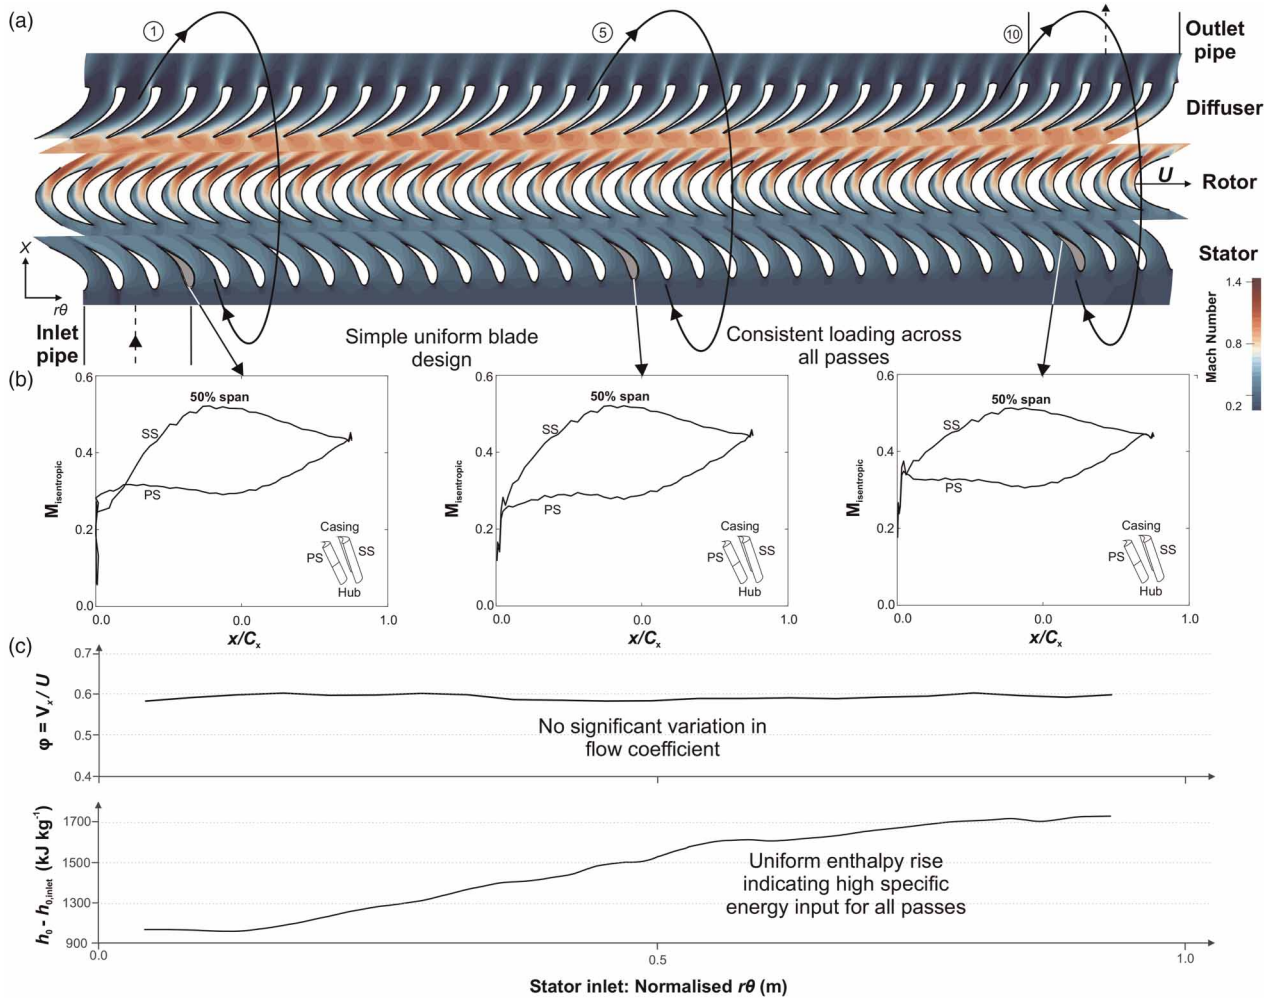


Figure 10. Time-averaged (URANS): (a) Mach number at mid-span across a linearised turbo-reactor view (b) isentropic Mach number $M_{isentropic}$ distribution for three stator passages across the turbo-reactor (c) flow coefficient ($\phi = (V_x/U)$) and total enthalpy $h_{r,0}$ referenced relative to the inlet conditions plotted circumferentially at the inlet to the stator.

Improving process quality: potential to reduce coke formation and deposition

Figure 11 illustrates that the new turbo-reactor provides the necessary conditions to potentially reduce coke formation and deposition. This is achieved through an order of magnitude higher average surface shear stresses

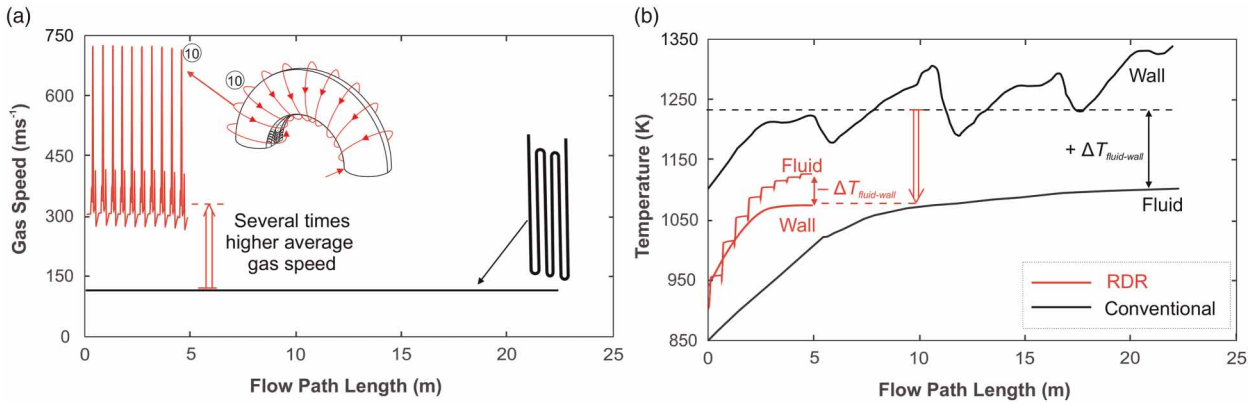


Figure 11. Factors affecting coke deposition in conventional reaction tubes (adapted from Gao et al. (2009); Niaei et al. (2002)) and the turbo-reactor: (a) gas speed (b) wall temperature and fluid-wall temperature gradient.

(due to several times higher throughflow speeds, see Figure 11a), a lower wall metal temperature (by ~200°C shown in Figure 11) and fluid–wall temperature gradient, in addition to 10× shorter residence times.

In sate-of-the-art conventional technology, a 60% reduction in the coking rate can be achieved by decreasing the tube metal temperature by 60°C (Vandewalle et al., 2017). This is considered a significant technological improvement. Therefore, a 200°C lower wall metal temperature in the RDR could lead to an order of magnitude reduction in the rate of coke deposition, indicating a potential step-change advancement in process quality. Consequently, this would lead to an increase in the plant availability and turbo-reactor lifetime, in addition to lower operating and maintenance costs. In contrast with 3-D internally ribbed radiant coil technologies in conventional reaction tubes, the improved coking rate is accomplished whilst avoiding large pressure losses, therefore mitigating any major losses to the olefin selectivities.

Controllability: tailoring the specific work input during operation

During plant operation, the level of specific enthalpy imparted into the working fluid needs to be tailored in order to maximise the yield (see Figure 1) depending on the real-time reaction dynamics and *a priori* modelling uncertainties. Furthermore, control of the enthalpy input and temperature distributions is vital in order to account for feedstock variability and flexibility (e.g., ethane, naphtha, butane, propane, bio-feeds, etc.), as well as to tailor the relative mass fractions of the products in the output stream depending on market price, customer requirements, and different downstream chemical processes.

The turbo-reactor offers the operational flexibility and controllability to produce relatively large changes in work input by performing only minor perturbations to the rotational speed. Figure 12 shows the envelopes of rotor blade loading and enthalpy rise across the machine resulting from a rotational speed variation around the nominal design value (Ω_{design}) of $\Omega_{\text{design}} \pm 5\%$. All other operational parameters, such as back pressure, are fixed. Nonetheless, during operation the back pressure can be controlled (within a small range) using a valve to further optimise the capacity and velocity triangles.

Figure 12 indicates that each stage successfully responds (almost linearly) to the change in rotational speed. Compared with non-dimensional similarity analysis i.e., $\Delta h_0 \propto \Omega^2$, the numerical results predict a similar envelope of stagnation enthalpy change across the turbo-reactor around the design enthalpy rise. The numerical results indicate that the enthalpy rise across the machine relative to the inlet reactor enthalpy ($\Delta h_{0,r}$ in Figure 12b) changes by ±7%, which is very similar to the ±10% change in enthalpy rise predicted by the similarity relation ($\Delta h_0 \propto \Omega^2$). At these off-design rotational speeds, deviations from this idealised analysis are primarily promoted by changes in incidence (due to the distortion of the design velocity triangles induced by insufficient mass flow changes), blockage, and separation levels, as well as lateral mixing. This is characteristic of a high-speed machine near the choke point since the change in capacity (mass flow rate) relative to the nominal value is sub-linear; in this case the mass flow rate changes by only ±3%.

Figure 13 summarises how the new turbo-reactor addresses many of the inherent limitations of a conventional furnace using an efficient and highly compact means of energy input into the working fluid.

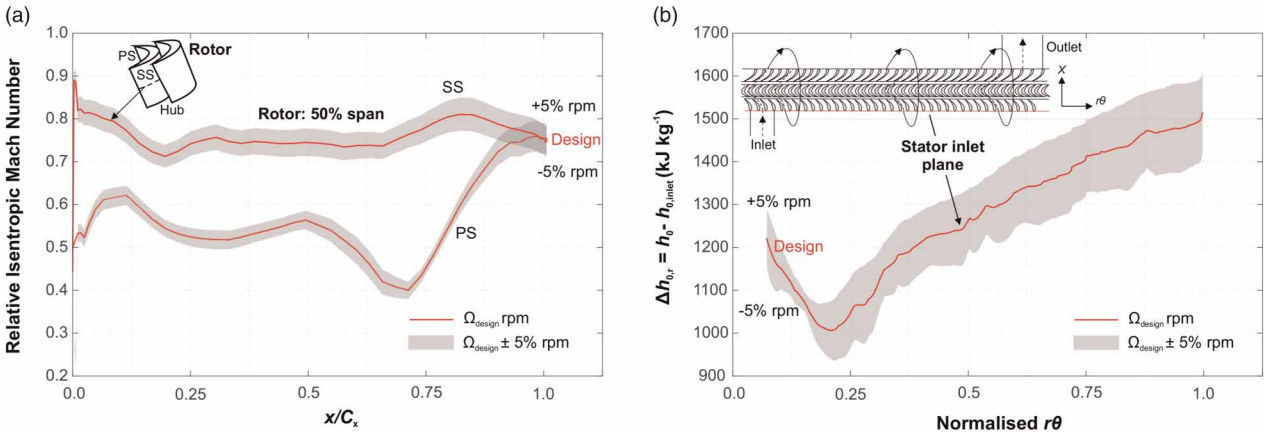


Figure 12. Operational envelopes (URANS) resulting from rotational speed perturbations $\Omega_{\text{design}} \pm 5\%$ about the design point: (a) rotor relative isentropic Mach number (b) change in stagnation enthalpy $\Delta h_{0,r}$ referenced relative to the reactor inlet enthalpy $h_{0,\text{inlet}}$, plotted upstream of the stator from reactor inlet to outlet.

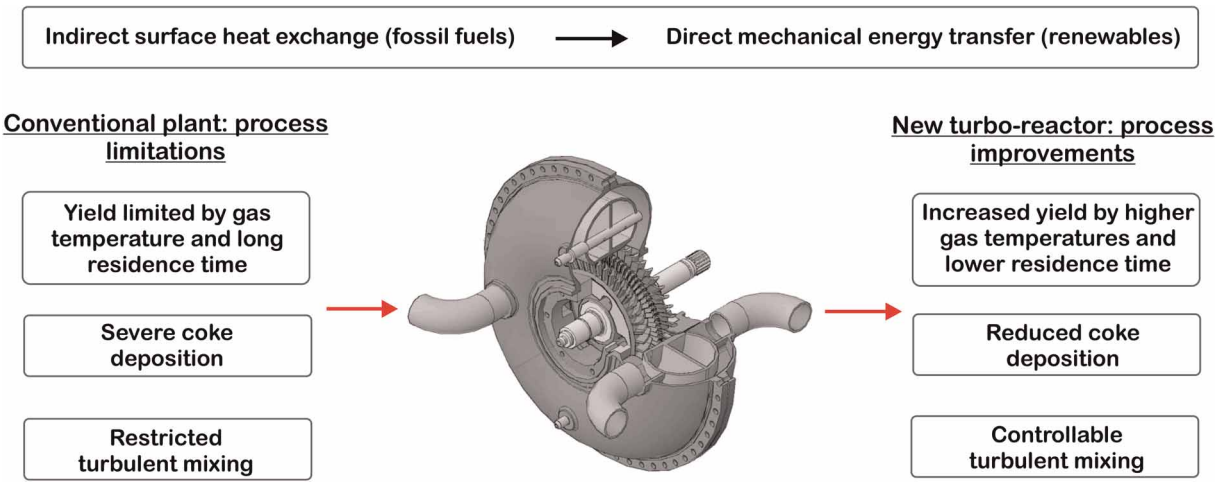


Figure 13. Block diagram summarising how the fundamental limitations of the conventional process are addressed.

Conclusions

1. This paper has introduced a novel concept aiming to replace the radiant section of a hydrocarbon cracking plant with a new turbo-reactor. The root of all the advantages of the RDR is in the replacement of the energy-input mechanism for the endothermic reaction from (indirect) surface heat exchange to direct mechanical energy transfer through a set of rotating blades. This paves the way to potentially decarbonise a highly energy-intensive, high-temperature industrial process by direct renewable electrification, as well as providing substantial improvements to the quality of the overall cracking process.
2. The turbo-reactor has the potential to remove the most fundamental limitations of the process, enabling a controllable and selective mechanism to generate the chemical reactions. Firstly, the residence time-scale for this process is up to an order of magnitude shorter than that of a conventional furnace, which provides a substantial improvement to the product yield and selectivity. This is enabled by supersonic throughflow speeds and impulse rotor blading, therefore enabling short time- and length-scale energy transfer and transformation processes. This includes an ultra-high loading rotor and a shock system in the diffuser. Secondly, a lower average wall temperature and a higher surface shear stress provide the necessary environment to potentially reduce coke formation and deposition, which is one of the most damaging side-effects of conventional cracking.
3. The numerical study presented has shown that the working principals of the elemental stage can be robustly realised in the full multi-stage environment. The necessary enthalpy input into the feedstock for the reaction can be provided regeneratively using only a single rotor with a short shaft, resulting a compact, low-cost reactor 500 times smaller than a conventional radiant furnace.
4. The robust unique open-volume vaneless space allows the rotor loading for each stage to be naturally de-sensitised to density variations, despite the application of a simple, uniform blade profile and inner/outer casing design (i.e., the distance between hub and casing is almost constant).
5. Intercommunication between the streamtubes, mediated by pressure waves propagating within the open-volume, guarantees an almost constant pressure equilibrium is maintained across the reactor as is required by the reaction. In conjunction with this, numerical simulations indicate that the intense turbulent mixing process is sufficient to further suppress a build-up of pressure within the non-bladed reactor space. Downstream of the mixing zone, the Reynolds number is very low which allows a rapid recovery of the flow to a well-conditioned state.
6. In spite of the challenges of the high relative Mach number design, it has been demonstrated that the turbo-reactor is able to respond to variations in rotational speed almost linearly and without deviating radically from the simple and ideal similarity rule. This indicates that sufficient speed-controllability is possible during operation. As a result, this increases the potential to perform online yield-tailoring to account for uncertainties in the reaction dynamics and feedstock variability and flexibility, as well as tailor the product yields depending on market price and customer requirements. This can be achieved by exploiting relatively large changes in specific energy imparted into the fluid resulting from only minor rotational speed perturbations.

Nomenclature

Symbols

C_x axial chord length, m
E_k specific kinetic energy, J kg^{-1}
h_0 stagnation enthalpy, J kg^{-1}

$$\mathrm{M}_{\mathrm{isentropic}} = \sqrt{\left(\left(\frac{p_{0,\mathrm{rel,inlet}}}{p}\right)^{\frac{\gamma-1}{\gamma}} - 1\right)\frac{2}{\gamma - 1}} \quad \text{isentropic Mach number}$$

p static pressure, Pa
p_0 stagnation pressure, Pa
T_0 stagnation temperature, K
U blade tip speed, m s^{-1}
u internal energy, J kg^{-1}
V_x axial velocity, m s^{-1}

$$\phi = \frac{V_x}{U} \quad \text{flow coefficient}$$

$$\psi = \frac{\Delta h_0}{U^2} \quad \text{work/loading coefficient}$$

Acronyms

RDR Roto-Dynamic Reactor
TE trailing edge

Acknowledgments

The authors would like to thank Coolbrook Oy for supporting this work and permitting the publication of this study. The authors also acknowledge the use of the University of Oxford Advanced Research Computing facility for this work.

Funding sources

Coolbrook Oy.

Competing interests

Dylan Rubini declares that he has no conflicts of interest. Nikolas Karefyllidis declares that he has no conflicts of interest. Liping Xu declares that he has no conflicts of interest. Budimir Rosic declares that he has no conflicts of interest. Harri Johannesdahl declares that he has no conflicts of interest.

References

Amghizar I., Dedeyne J. N., Brown D. J., Marin G. B., and Van Geem K. M. (2020). Sustainable innovations in steam cracking: CO_2 neutral olefin production. Reaction Chemistry & Engineering. 5 (2): 239–257. https://doi.org/10.1039/C9RE00398C

Bushuev V. A. (2016). Bladed reactor for the pyrolysis of hydrocarbons . U.S. Patent Application No. US 9.494,038 B2. Washington, DC: U.S. Patent and Trademark Office.

Chauvel A. and Lefebvre G. (1989). Petrochemical processes (volume 1: Synthesis-gas derivatives and major hydrocarbons), Vol. 1. Editions Technips. Paris, France.

Coolbrook Oy (2020), https://coolbrook.com. Accessed: 15/10/2020.

de Pee A., Pinner D., Roelofsen O., Somers K., Speelman E., and Witteveen M. (2018). Decarbonization of industrial sectors: the next frontier. Amsterdam, The Netherlands: McKinsey & Company.

Denton J. D. (1982). An improved time marching method for turbomachinery flow calculation, Vol. 1: Turbomachinery of Turbo Expo: Power for Land, Sea, and Air. V001T01A082.

Denton J. D. (1986). The use of a distributed body force to simulate viscous effects in 3D flow calculations, Vol. 1: Turbomachinery of Turbo Expo: Power for Land, Sea, and Air. V001T01A058.

Denton J. D. (1992). The calculation of three-dimensional viscous flow through multistage turbomachines. Journal of Turbomachinery. 114 (1): 18–26. https://doi.org/10.1115/1.2927983

Gao G.-Y., Wang M., Ramshaw C., Li X.-G., and Yeung H. (2009). Optimal operation of tubular reactors for naphtha cracking by numerical simulation. Asia-Pacific Journal of Chemical Engineering. 4 (6): 885–892. https://doi.org/10.1002/apj.351

Haydock H., and Napp T. (2013). Decarbonisation of heat in industry: A review of the research evidence, Ricardo-AEA and Imperial College Consultants.

IRENA (2014). Renewable Power Generation Costs in 2014, Technical report, The International Renewable Energy Agency.

IRENA (2016). Renewable Capacity Statistics 2016, Technical report, The International Renewable Energy Agency, Abu Dhabi.

Jambor B. and Hájeková E. (2015). Formation of coke deposits and coke inhibition methods during steam cracking. Petroleum & Coal. 57 (2): 143–153.

Keller F., Lee R. P., and Meyer B. (2020). Life cycle assessment of global warming potential, resource depletion and acidification potential of fossil, renewable and secondary feedstock for olefin production in germany. Journal of Cleaner Production. 250: 119484. https://doi.org/10.1016/j.jclepro.2019.119484

Klostermeier C. (2008). Investigation into the capability of large eddy simulation for turbomachinery design. PhD thesis, University of Cambridge.

Linde Engineering (2020). Linde industrial steam cracking furnace model, https://www.linde-engineering.com/en/index.html. Accessed: 15/10/2020.

Mynko O., Amghizar I., Brown D. J., Chen L., Marin G. B., et al. (2022). Reducing CO2 emissions of existing ethylene plants: evaluation of different revamp strategies to reduce global CO2 emission by 100 million tonnes. Journal of Cleaner Production. 362: 132127. https://doi.org/10.1016/j.jclepro.2022.132127

Niaei A., Towfighi J., Sadrameli M., and Masoumi M. E. (2002). Computational study of the pyrolysis reactions and coke deposition in industrial naphtha cracking. In International Conference on Computational Science. Springer.

Rosic B., Denton J. D., and Pullan G. (2005). The importance of shroud leakage modeling in multistage turbine flow calculations. Journal of Turbomachinery. 128 (4): 699–707. https://doi.org/10.1115/1.2181999

Rubini D., Xu L., Rosic B., and Johannesdahl H. (2021). A new turbomachine for clean and sustainable hydrocarbon cracking. Journal of Engineering for Gas Turbines and Power. 144 (2): 021024. https://doi.org/10.1115/1.4052784

Schietekat C. M., Van Cauwenberge D. J., Van Geem K. M., and Marin G. B. (2014). Computational fluid dynamics-based design of finned steam cracking reactors. AIChE Journal. 60 (2): 794–808. https://doi.org/10.1002/aic.14326

Sixsmith H. and Altmann H. (1977). A regenerative compressor. Journal of Engineering for Industry. 99 (3): 637–647. https://doi.org/10.1115/1.3439291

Van Goethem M. W. M. (2010). Next generation steam cracking reactor concept. PhD thesis, Delft University of Technology.

Vandewalle L. A., Van Cauwenberge D. J., Dedeyne J. N., Van Geem K. M., and Marin G. B. (2017). Dynamic simulation of fouling in steam cracking reactors using CFD. Chemical Engineering Journal. 329: 77–87. https://doi.org/10.1016/j.cej.2017.06.113

Yuan B., Zhang Y., Du W., Wang M., and Qian F. (2019). Assessment of energy saving potential of an industrial ethylene cracking furnace using advanced exergy analysis. Applied Energy. 254: 113583. https://doi.org/10.1016/j.apenergy.2019.113583

Zhang J., Van de Vijver R., Amghizar I., Reyniers M.-F., Van Geem K. M. (2022). Combined Catalytic and Pyrolytic Coking Model for Steam Cracking of Hydrocarbons. Industrial & Engineering Chemistry Research. 61 (11): 3917–3927. https://doi.org/10.1021/acs.iecr.1c05036

Statement of Authorship for joint/multi-authored papers for PGR thesis

To appear at the end of each thesis chapter submitted as an article/paper

The statement shall describe the candidate's and co-authors' independent research contributions in the thesis publications. For each publication there should exist a complete statement that is to be filled out and signed by the candidate and supervisor **(only required where there isn't already a statement of contribution within the paper itself).**

Title of Paper	A New Robust Regenerative Turbo-reactor Concept for Clean Hydrocarbon Cracking
Publication Status	☐Published ☐ Accepted for Publication ☐Submitted for Publication ☐Unpublished and unsubmitted work written in a manuscript style
Publication Details	Rubini, Dylan, Nikolas Karefyllidis, Liping Xu, Budimir Rosic, and Harri Johannesdahl. "A new robust regenerative turbo-reactor concept for clean hydrocarbon cracking". Journal of the Global Power and Propulsion Society 6 (2022): 135-150. doi:10.33737/jgpps/150550.

Student Confirmation

Student Name:	Dylan Rubini		
Contribution to the Paper	Performed some numerical simulations, developed required codes for simulation and postprocessing, post-processed and analyzed results and wrote manuscript		
Signature		Date	16/08/2024

Supervisor Confirmation

By signing the Statement of Authorship, you are certifying that the candidate made a substantial contribution to the publication, and that the description described above is accurate.

Supervisor name and title: Prof. Budimir Rosic			
Supervisor comments I confirm that the candidate's statement is an accurate description of his contribution to this publication.			
Signature		Date	19/08/2024

This completed form should be included in the thesis, at the end of the relevant chapter.

2.2.2 Perspective

It has been understood that, because of constraints imposed by the heat transfer mechanism of a furnace, the pyrolysis efficiency appears to have reached its practical maximum for conventional designs. This is a consequence of the fluid temperature level being limited by metallurgical constraints, and the residence time being limited by the requirement of having a large surface heat exchange area ($\dot{Q} = U_h A \Delta T$).

In the early 1990s, the SWR was developed to address some of these challenges and improve the reaction yield. However, this technology lacked controllability and was not carbon neutral. These challenges are overcome by the turbo-reactor. Rather than approximating volumetric energy transfer as in the SWR, the turbo-reactor provides true direct mechanical energy transfer into the fluid at a much lower peak *system* temperature level than conventional or SWR technologies. As a result, higher *gas* mixture temperatures are feasible, and thus substantial improvements in reaction performance.

The fundamental elemental stage design has been introduced, and the effectiveness of an ultra-high stage loading rotor row coupled with a rapid dissipation process has been demonstrated. Turbulent mixing between stages redistributes the chemical species and temperature, promoting reaction uniformity. This level of uniformity is not possible in conventional technology. Despite this upstream turbulent mixing, the results of high-fidelity numerical simulations indicate that the flow recovers sufficiently between passes, ensuring high loading conditions for each regenerative pass and confirming the feasibility of a regenerative architecture. Numerical calculations also confirm that the streamtubes naturally self-adjust to promote fairly uniform conditions entering successive passes.

Why (Not) Regenerative?

The regenerative concept provides four crucial advantages. First, the shaft length is short, which benefits the rotordynamics and thermomechanics. Second, only a single rotor blade row is needed to achieve multi-stage work input equivalent to 10 to 12 stages, which can reduce capital costs. Third, since the rotor row rotates through hot and cold temperature zones, the mean effective temperature "seen" by each rotor passage is approximately an average of the cold inlet and hot outlet temperatures. This self-induced cooling effect by

the working fluid itself eliminates the need for expensive and intricate blade cooling systems and is a remarkable benefit of the regenerative design. This could lead to an improved service life of the rotor and potentially almost eliminate coking within the rotating blade row. Finally, the regenerative design allows natural streamtube adjustment and pressure equilibration in the open volume. As a result, despite streamwise density variation, the flowfield entering each successive pass is fairly constant without needing a non-uniform blading or casing design. The same cannot necessarily be said for an axial architecture.

Despite the advantages of the regenerative design, it is not without challenges compared to other configurations. These can be summarised as follows.

(1) Capacity is limited for the regenerative design and is a main motivation for an axial architecture.

(2) Although axial thermal stresses are minimised, they are now replaced by circumferential thermal stresses around the casing. Potentially, these could be mitigated by adopting a thin flexible casing structure, which is a plausible solution since the pressure differential across the casing is only 1 bar.

(3) Since the regenerative passes are not segregated, heat and mass can transport across pass boundaries. This means that pyrolysed and un-pyrolysed components of the gas mixture can mix, which will reduce the feed reaction conversion in the effluent gas stream. However, this can be solved by segregating the passes.

(4) Although the rotor "sees" an effective temperature that is the *average* of the cold inlet and hot outlet, the creep life is a non-linear function of temperature. Therefore, the *peak* rotor temperature will be limiting. Although this temperature disturbance will not penetrate deep below the metal ($\delta_p \sim \text{constant} \cdot \sqrt{\frac{\alpha_T}{f}} < 0.5\,\text{mm}$) benefiting bulk creep life, low-frequency thermal fatigue at the blade surface may lead to surface cracks.

(5) In the regenerative architecture, the temperature profile is not controllable since the multi-stage design is not modular. Chapter 3 will show that instead the modularity and controllability of an *axial* multi-stage design enables temperature profile tailoring for a more efficient chemical reaction.

Throughout the remainder of this thesis, we will focus on an axial architecture, which is better suited for large-scale capacity and has the potential for important reaction-guided temperature profiling.

Chapter 3

Redefining Energy Quality for the Turbo-Reactor

3.1 Energy Conversion Mechanisms

In the preceding chapter, a fundamental understanding of the working principles of the elemental stage and its aerodynamics has been developed. Without the requirement for a pressure rise, efficiency and stability restrictions are removed, enabling ultra-high specific enthalpy rates and power densities. To facilitate improvements in reaction selectivity, rapid mechanical energy input must be complemented with a controllable dissipation process. To achieve this, there is an important change in design philosophy. Fundamentally, conventional turbomachinery design wisdom tells us to minimise dissipation (*i.e.,* entropy generation or aerodynamic losses), whereas, in the turbo-reactor, the objective is to controllably increase entropy production to "design" a reaction-efficient temperature profile. To understand this, it will be explained that reaction efficiency & selectivity and chemical exergy must be included as new vectors of energy quality for turbomachines. As a result, this chapter aims to reframe the turbo reactor as a "chemical-reaction-tuned" energy conversion device. The modular and customisable multi-stage axial architecture, as opposed to the regenerative environment, is better suited to optimising the temperature profile for ideal reaction efficiency.

Because the species net production/destruction rates $\dot{\omega}_k$ (*i.e.,* the reaction dynamics) depend exponentially on the local static temperature level T through the Arrhenius equation

$$\dot{\omega}_k(T,\ p,\ Y_1, \ldots, Y_K) \propto A_f T^{\beta} \exp\left(\frac{-E_a}{R_u T}\right) \tag{3.1}$$

temperature is the most important thermodynamic variable that influences the chemistry. In contrast, the reaction dynamics typically exhibit only a first-order dependence on pressure ($\dot{\omega}_k(T,\ p,\ Y_1, ..., Y_K) \propto p$), so this is less important.

Two control mechanism can be used to influence the energy absorption rate by the reaction:

1. *Coarse-grained control:* adjusting the interstage vaneless space length

2. *Fine-grained control:* adjusting the local entropy generation rate

This chapter quantifies the breakdown of physical mechanisms responsible for kinetic to internal energy conversion. In future design iterations, an improved understanding of the physics will allow better control of the energy supply rate to the chemical reaction. As a result, careful aerodynamic tailoring of the bladed path can be employed to design an ideal temperature profile for the reaction.

Pages 55–68:

Energy Conversion Mechanisms in Turbomachines for High-Temperature Endothermic Reactions: Redefining Energy Quality

(*Appl. Therm. Eng.* 2025)

Applied Thermal Engineering 258 (2025) 124566

Contents lists available at ScienceDirect

Applied Thermal Engineering

journal homepage: www.elsevier.com/locate/ate

Research Paper

Energy conversion mechanisms in turbomachines for high-temperature endothermic reactions: Redefining energy quality

Dylan Rubini [a,*], Budimir Rosic [a], Liping Xu [b]

[a] Oxford Thermofluids Institute, Department of Engineering Science, University of Oxford, Oxford OX2 0ES, UK
[b] Whittle Laboratory, Department of Engineering, University of Cambridge, Cambridge CB3 0DY, UK

ARTICLE INFO

Keywords:
Turbomachinery
Endothermic chemical processes
Steam cracking
Turbo-reactor
Loss breakdown

ABSTRACT

A new class of turbomachines, called turbo-reactors, have emerged to decarbonize high-temperature chemical processes. These applications unlock a new aerothermochemical design space for turbomachines. This paper explains how the turbo-reactor has the potential to be a "chemically tuned" device. In contradiction with conventional design wisdom, this can be achieved by balancing entropy generation within the flow against heat absorption by the reaction. This enables the "design" of a reaction-efficient temperature profile. To do this, it is necessary to understand and quantify the distribution of the entropic and isentropic mechanisms responsible for energy conversion. This paper uses high- and low-fidelity computations to decompose the energy conversion process into mechanisms based on a spatial decomposition. A range of Mach and Reynolds number regimes are studied, as well as a multistage configuration without vaneless space. The energy conversion breakdown analysis indicates that the entropic energy conversion dominates over the isentropic component with a contribution of 65%. The dominant source of entropy production is viscous dissipation generated by the thick diffuser trailing edge, accounting for 25% of the total. The shock system provides 20% of the energy conversion, almost entirely due to reversible pressure rise rather than entropy production. The energy conversion coefficient is independent of Reynolds number over engine-relevant conditions, whereas Mach effects are more significant. Across the Mach numbers range 1.1 to 1.5, the energy conversion coefficient increases by 20%. This is lower than expected as a result of the opposing effects of reversible and irreversible energy conversion contributions.

1. Introduction

1.1. Historical perspective, motivation and background

The fundamental principles of turbomachinery design have been shaped to minimize the degradation of energy quality, *i.e.,* the exergy destruction or lost work potential, during energy conversion processes in power generation, aviation and industrial processes [1,2]. The focus on reducing energy quality degradation (and increasing fuel efficiency) intensified following the 1970s oil crisis, reaffirming that the finite fossil fuel resource was rapidly depleting, and encouraging industries to strive towards extracting the maximum utilizable work, that is, the exergy, from nonrenewable energy resources. Today, the drive to minimize aerodynamic losses for efficiency gains is ubiquitous in turbomachinery [3].

In the context of conventional compressors (*i.e.,* energy-imparting machines), the bladed flow path is designed to maintain the absorbed shaft power in the form of pressure or kinetic energy to preserve energy quality. However, now with the emergence of new applications for turbomachines — decarbonizing high-temperature endothermic reaction processes with a new turbo-reactor [4,5] — this is no longer the case. Specifically, the goal of the turbo-reactor seen in Fig. 1 is to convert energy from the electric-motor-powered shaft into internal energy of the fluid (not pressure). Since there is no global pressure rise, this can be achieved effectively and without instability using ultrahigh-loading rotating blades to impart mechanical energy to the fluid before it is dissipated (and converted) into internal energy. The working fluid undergoes a chemical reaction along most of the bladed flow path, which presents a new set of requirements for the design. That is, the imparted mechanical energy must be dissipated at a rate determined by the reaction, which is typically relatively fast. This means that enhanced entropy generation is not only unavoidable but necessary. This new turbomachinery design principle of leveraging entropy generation for ultrafast energy transformation along the bladed path is fundamentally different from conventional wisdom.

1.2. Objectives of this paper

Following a recap of the basic requirements and key features of the turbo-reactor concept that has been introduced in detail in previous

* Corresponding author.
E-mail address: dylan.rubini@eng.ox.ac.uk (D. Rubini).

https://doi.org/10.1016/j.applthermaleng.2024.124566
Received 17 June 2024; Received in revised form 12 September 2024; Accepted 4 October 2024
Available online 16 October 2024
1359-4311/© 2024 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Nomenclature

Roman Letters

b_{true}	True diffuser chord length, m
c_p	Isobaric heat capacity, $\text{J kg}^{-1}\text{ K}^{-1}$
C_x	Diffuser axial chord length, m
h_0	Total enthalpy, J kg^{-1}
p	Static pressure, bar
t_{TE}	Diffuser trailing edge thickness, m
T	Static temperature, K
u_i	Velocity vector, m s^{-1}
U	Mean blade speed, m s^{-1}
V_1	Diffuser inlet velocity, m s^{-1}
w	Diffuser blade pitch, m

Greek Letters

ζ_h	Total energy conversion coefficient
$\zeta_{h,\text{rev.}}$	Reversible (isentropic) energy conversion coefficient
$\zeta_{h,s}$	Irreversible (entropic) energy conversion coefficient
μ	Dynamic molecular viscosity, Pa s
ρ	Density, kg m^{-3}
$\bar{\tau}_{ij}$	Mean flow shear stress tensor, Pa
τ^R_{ij}	Reynolds stress tensor, Pa

Dimensionless Groups

M	Mach number
Re	Reynolds number
σ	Shock function, $\frac{\boldsymbol{M}\cdot\nabla p}{\lvert\nabla p\rvert}$

Superscripts

$\square'$	Fluctuating component
$\overline{\square}$	Time averaged

Acronyms

BL	Boundary layer
ITV	Integrated turning vane
LES	Large-eddy simulation
LE/TE	Leading/trailing edge
PS/SS	Pressure/suction surface
SWBLI	Shockwave boundary layer interaction
URANS	Unsteady Reynolds–averaged Navier–Stokes

work by the authors [4–7], the paper is organized into four distinct parts summarized below.

(1) The reaction requirements are outlined along with how the turbo-reactor is designed to achieve them.
(2) The consequences of converting electricity to heat and thus degrading energy quality are clarified. By introducing chemical exergy as a new vector of energy quality for turbomachines, it will be demonstrated that this process is sensible from a thermodynamic perspective.
(3) This paper demonstrates how the turbo-reactor can be designed as a "chemically reaction tuned" — or for brevity, a "chemically tuned" — device (and is not a "bad compressor" or a "mixer"). It is explained how aerodynamic losses can be exploited to design a reaction-efficient temperature profile by controlling the dissipation profile. This unlocks a new dimension within the design space of turbomachines.
(4) To design this reaction-efficient temperature profile, an improved understanding of the energy conversion mechanisms is required. Therefore, the distribution of energy conversion mechanisms is numerically investigated over a range of design conditions. An improved understanding of this process will inform future aerodynamic design optimization to tune the energy conversion process for increased reaction efficiency.

The focus of this paper is on reacting flow applications. For context, the example endothermic reaction application of steam cracking of hydrocarbon feedstocks [4,5] is used. However, Rubini et al. [7] demonstrated that the turbo-reactor can also be used to decarbonize a variety of other endothermic reaction processes. The reader is referred to Chauhan et al. [8], Mynko et al. [9], and Thiel & Stark [10] for a comprehensive review of alternative pathways to decarbonize high-temperature chemical reaction processes, which include electrochemical pathways, alternative combustion fuels such as hydrogen, plasma pyrolysis, and electric furnaces using resistance (Joule) heating.

2. Decarbonizing endothermic reactions: The turbo-reactor concept

The renewably-powered turbo-reactor replaces fossil-fuel-fired tubular furnaces. For an ethane feedstock, Fig. 2(*a*) compares a typical turbo-reactor stage with a small section of a furnace pipe of equivalent length. In the furnace tube, the rate of energy transfer is limited by the thermal diffusion timescale of surface heat exchange. Moving to direct volumetric energy input within the turbomachine (see Fig. 2), the energy transfer rate is instead driven by the rotor convection timescale. Whereas the thermal diffusion timescale is essentially fixed, the convection timescale can be compressed through high supersonic flow speeds, as is done in the turbo-reactor. Therefore, orders of magnitude higher power densities can be achieved relative to conventional furnaces (see Fig. 2(*a*) & (*e*)).

2.1. Reaction requirements

For turbo-reactor-powered endothermic chemical reaction processes, value is created by chemically transforming low-value feedstock (*e.g.* ethane) into higher-value products (*e.g.,* ethylene). Alongside power input and power density targets, the efficiency of this chemical transformation — that is, the fraction of feedstock that is converted into high-value product species — is the most important figure of merit and displaces isentropic stage efficiency as the driving design metric for this new class of machines. Achieving a highly efficient chemical reaction requires two design features, both of which are now possible by replacing furnaces with the turbo-reactor.

(1) The reaction should progress uniformly in the pitchwise and spanwise directions. The progress of the reaction is most significantly impacted by the temperature field. In the pipes of conventional furnaces, Fig. 2(*c*) illustrates that the temperature distribution is highly nonuniform due to heat transfer through the walls. This results in reaction progression along each streamline going out of sync, as shown by the radial ethane feed conversion difference ΔY_C highlighted in Fig. 2(*d*). This leads to poor reaction efficiency and solid carbon (*i.e.,* coke) forming directly on the metal surface. Fundamentally, this challenge of nonuniformity cannot be avoided even with novel electric or hydrogen furnaces [11], which are still surface heat exchangers. In contrast, in the turbo-reactor, the volumetric energy input mechanism combined with turbulent mixing between stages facilitates a significantly more uniform temperature distribution (see Fig. 2(*b*)). This can promote higher reaction efficiency.

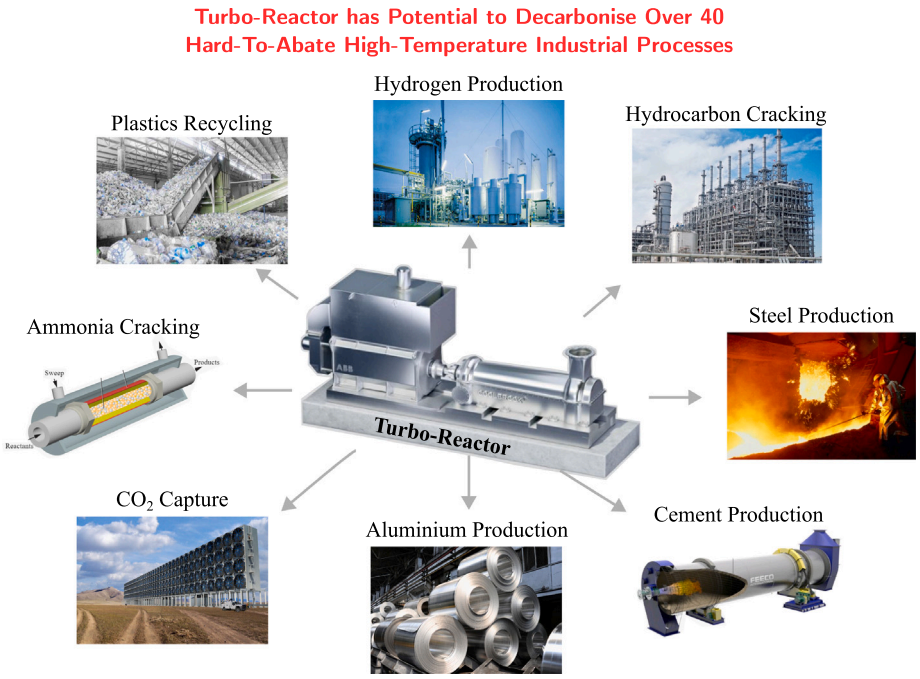


Fig. 1. A turbo-reactor render showing some example industries that it could decarbonize.

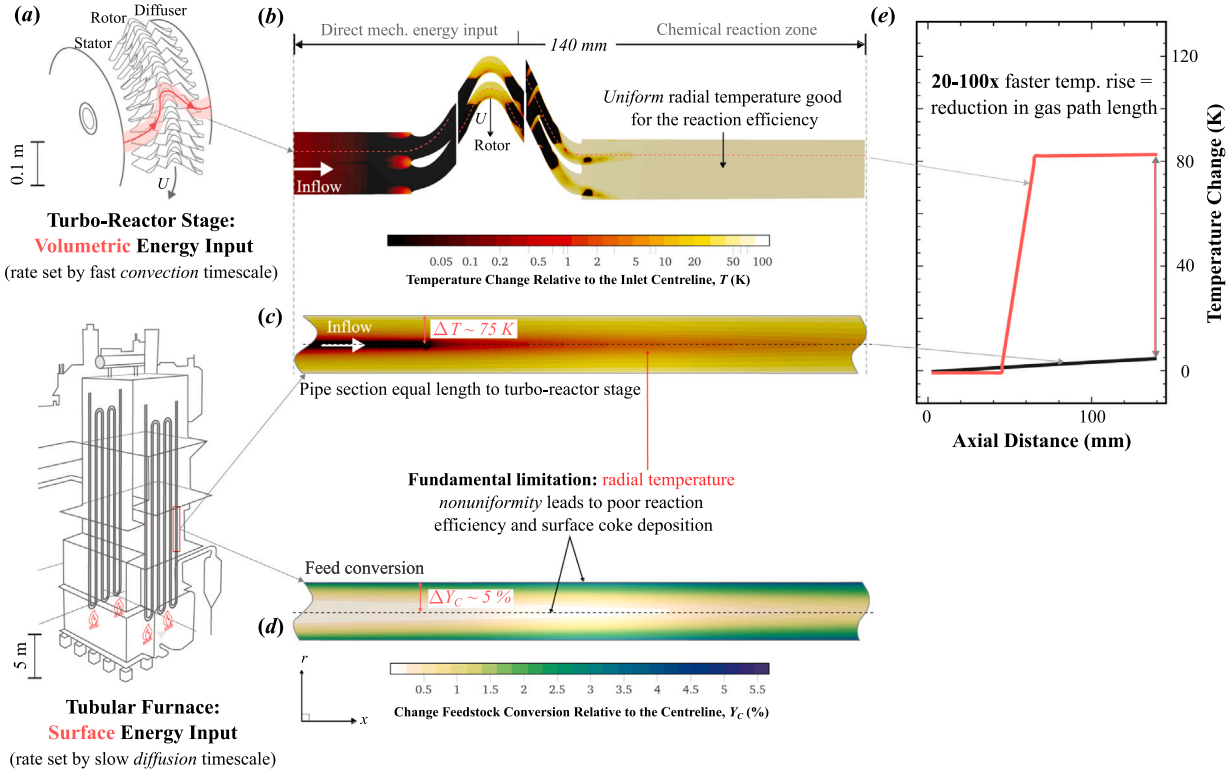


Fig. 2. (*a*) Schematic comparison between a turbo-reactor stage and a typical steam cracking furnace pipe (*b*) & (*c*) a comparison of the static temperature change on a *logarithmic* scale (*d*) change in ethane feedstock conversion change and (*e*) a comparison of the averaged axial temperature change.

(2) To increase reaction efficiency, shorter reaction timescales at higher static temperature levels should be designed for, as shown in Fig. 3. Bypassing metallurgical limitations due to the hot walls of conventional surface heat exchanges allows the turbo-reactor to achieve higher bulk fluid temperatures over a shorter reaction timescale. This improvement in reaction efficiency is a key advantage of the turbo-reactor that is not possible with surface-heat-exchange-based methods to decarbonize high-temperature chemical processes (*e.g.,* hydrogen or electric furnaces) [11].

We will see that these features of the turbo-reactor design are enabled by controllably enhancing aerodynamic losses and viscous dissipation.

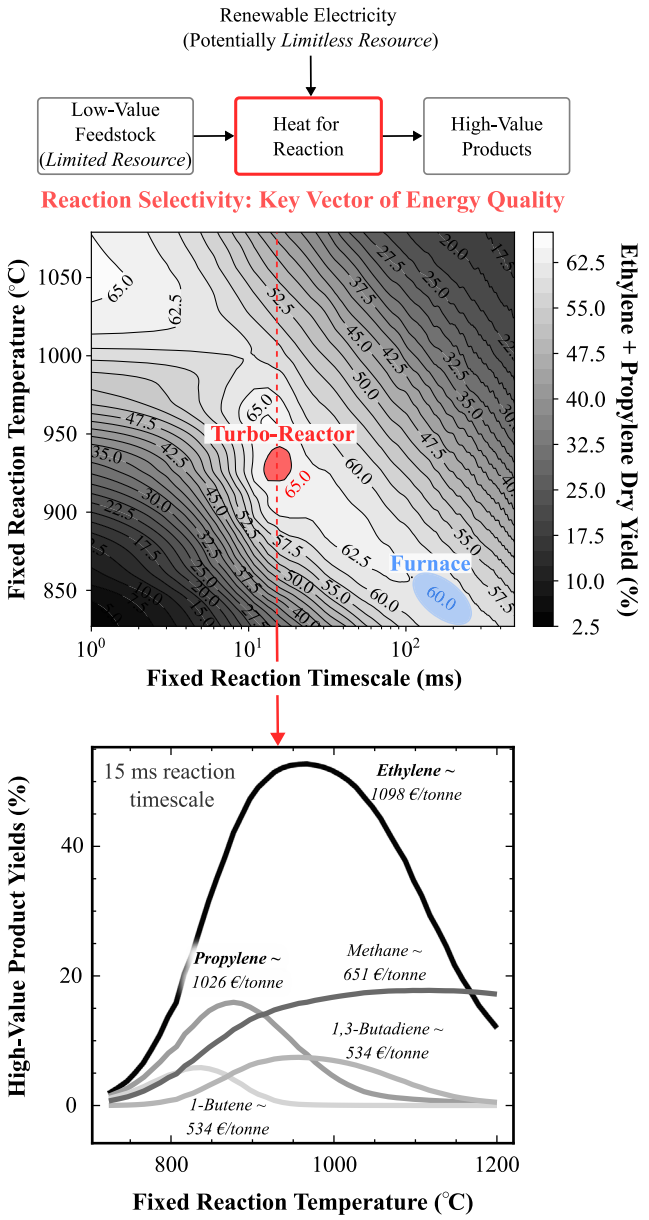


Fig. 3. Idealized dry yields of selected products as a function of fixed reaction temperature and timescale. The lower figure show the product yields as a function of temperature at a fixed reaction timescale 25 ms as highlighted by the red dashed line.

The next section summarizes the elemental stage design philosophy and approach to configuring the multistage environment to achieve the reaction requirements presented in this section.

2.2. Working principles of the turbo-reactor

2.2.1. Energy transfer

Short reaction timescale requirements dictate that an ultrahigh loading impulse rotor is employed. Since there is no requirement for a global pressure gain, it is possible to develop supersonic rotor blades with a work coefficient of up to $\psi = \frac{\Delta h_0}{U^2} \approx 8.0$, depending on the fluid properties [7]. This is 20 times that of an axial compressor. Without a global pressure rise, efficiency and stability restrictions that limit the work coefficient in compressors are alleviated [12], enabling this loading level.

2.2.2. Energy conversion

To supply an enthalpy dh to the chemical reaction at a constant pressure, the Gibbs equation

$$dh = \frac{1}{\rho} dp + T\,ds \tag{1}$$

implies that an entropy change ds must be produced to satisfy the constant global static pressure ratio ($dp \approx 0$) boundary conditions. Since the designer is no longer concerned with efficiency, aerodynamic losses (and hence entropy production and viscous dissipation) can be leveraged to control the energy conversion rate profile to design a temperature profile that favors optimal reaction dynamics (see Section 3.1).

2.2.3. Chemically tuning the interstage vaneless space

Fig. 4 shows that each stage is separated by an interstage vaneless mixing space where the mechanical energy absorbed through the rotor is converted into internal energy. This energy conversion process must be distributed over a series of stages to counterbalance the drop in temperature due to the endothermic chemical reaction, as shown in Fig. 5. In Fig. 3 it was suggested that the local residence time favored by the reaction varies with the temperature level, which is different for each stage. Therefore, the local residence time can be controlled by varying the length of the vaneless space between stages (L). This becomes a new tunable control parameter (see Fig. 4), which can be customized stage by stage from length 0 to NC_x to design a reaction-efficient temperature profile (see Fig. 5). Furthermore, this vaneless mixing space is also critical to allow thermochemical nonuniformities to mix-out. The design of the temperature profile for efficient chemistry is explained in more detail in Section 3.1.

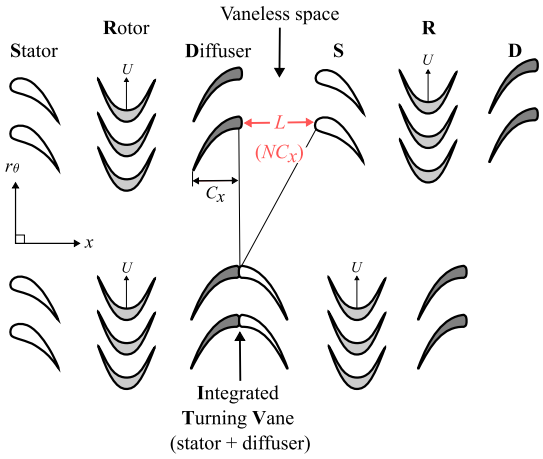


Fig. 4. Variable interstage vaneless space length from 0 diffuser chords (C_x), *i.e.,* an integrated turning vane, up to NC_x.

For preheating the gas mixture before reaction activation, there are no reaction timescale requirements, so a "compact stage" design can be used. This is achieved by combining the stator and diffuser into an integrated turning vane (ITV), as shown in Fig. 4.

2.3. New vector of energy quality (Exergy)

It is evident that the chemical processes driven by the turbo-reactor rely on *irreversibilities* to rapidly dissipate mechanical energy within the required rapid reaction timescales (see Fig. 3) while promoting reaction uniformity (and hence performance). Therefore, it is important to assess the thermodynamic implications of this degradation of energy quality (exergy).

Before doing so, a new vector of energy quality must be introduced for this new class of turbomachines. Conventionally, turbomachinery design practice concentrates only on *physical* exergy. However, in the new endothermic reaction applications, energy is supplied to a chemical reaction along the bladed flow path. This means that *chemical*

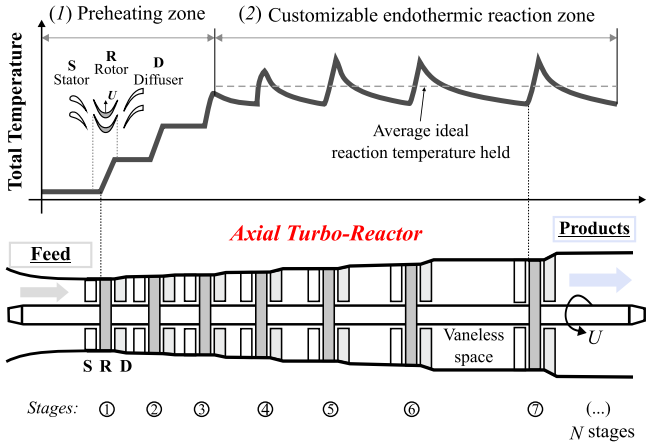


Fig. 5. Schematic of the elemental stage and multistage environment for the turbo-reactor, illustrating a typical total temperature profile for an endothermic chemical process.

exergy, defined in Appendix, must be considered in addition to physical exergy as a new vector of energy quality. This is an entirely new design consideration for turbomachines.

Equipped with chemical exergy, as well as the price of energy and the cost of feed/products, the impact of exploiting viscous dissipation can be understood using Fig. 6. The chemical exergies are calculated using the equations presented in Appendix. The inlet & output mass fraction distributions and reaction irreversibility are calculated from 1D plug-flow reactor chemical kinetic calculations. Aerodynamic irreversibility is determined by the predicted entropy generation from CFD calculations, and motor & heat losses are set to typical values for turbomachines. The energy price and the cost of the feedstock & products are estimated from the current (2024) average electricity prices and average hydrocarbon market prices [13], respectively. The energy requirement is set to that needed for the preheating and reaction completion.

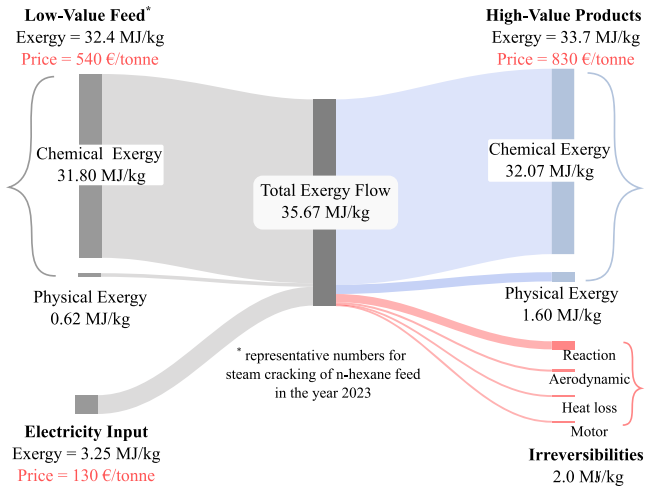


Fig. 6. Sankey diagram showing an exergy balance for a representative steam cracker with an n-hexane feed in 2023.

Fig. 6 shows that the supplied electricity has a significantly lower cost and exergy content compared to the feedstock. Therefore, based on the values in Fig. 6, combined with the assumption that the electricity supplied is semi-limitless (*i.e.,* provided by renewable energy), it is reasonable to sacrifice the energy quality supplied to the turbo-reactor by developing irreversibilities. This enables higher reaction efficiencies. After all, a higher reaction efficiency means more species in the effluent gas stream that have a high value. Therefore, it is useful to view reaction efficiency as a new vector of energy quality.

3. Leveraging aerodynamic losses for a controllable energy conversion profile

3.1. Chemically-tuned turbomachines: Design of a reaction-efficient temperature profile

In Section 2 it was highlighted that at the global *coarse-grained* level, the turbo-reactor is designed to operate at a higher temperature level over shorter timescales to improve reaction efficiency. We saw that the temperature could be controlled at the coarse-grained stage-level (see Section 2.2.3) by customizing the interstage vaneless space distance.

At the *fine-grained* level, there is not only an optimal reaction temperature level and timescale, but also an ideal temperature profile that benefits reaction performance. This section explains how the local energy conversion process, and thus the profile of energy absorbed by the reaction, could be tuned to "design" an ideal reaction-efficient temperature profile. The ability to perform temperature profiling is an important benefit of distributing energy addition across a series of stages. We will explore the potential benefits of tailoring the energy conversion profile at a fine-grained level, opening new possibilities for future aerothermochemical design optimization. The following discussion will be structured following the three boxes displayed in Fig. 7.

3.1.1. Box ①: Objectives

In an idealized scenario, the temperature level is raised and maintained (see Fig. 7) at an ideal reaction temperature set by the given residence time (see Fig. 3). To better understand how to design for this reaction temperature profile, it is useful to consider the temperature transport equation for the time-averaged base flow [14]. For an ideal gas in simplified form, this is defined as follows:

$$\bar{\rho}c_p\frac{D\bar{T}}{Dt} \equiv \bar{\rho}c_p\bar{u}_j\frac{\partial\bar{T}}{\partial x_j} \approx \underbrace{\lambda_{eff.}\frac{\partial^2\bar{T}}{\partial\bar{x}_j^2}}_{\text{diffusion}} + \Bigg(\underbrace{\bar{u}_i\frac{\partial\bar{p}}{\partial x_i}}_{\text{pressure work}} + \underbrace{(\bar{\tau}_{ij}+\tau_{ij}^R)\frac{\partial\bar{u}_i}{\partial x_j}}_{\text{viscous work}}\Bigg) - \underbrace{\dot{\omega}_T}_{\text{heat sink}} \tag{2}$$

where $\bar{\rho}$ is the density, c_p is the isobaric heat capacity, $\lambda_{\text{eff.}}$ is the effective thermal conductivity, $\bar{\tau}_{ij}$ is the stress tensor of the mean flow, τ_{ij}^R is the Reynolds stress tensor and u_i is the velocity vector. The terms in red represent different routes for the conversion of mean flow kinetic energy into enthalpy through pressure work and molecular & turbulent viscous shear work (*i.e.,* viscous dissipation). The term in blue $\dot{\omega}_T$ is the endothermic heat sink rate and is defined as

$$\dot{\omega}_T \approx \sum_{k=1}^{N_{\text{spe.}}} \Delta h_{f,k}^0 \cdot \dot{\omega}_k \tag{3}$$

where $\dot{\omega}_k$ are the net species production rates and $\Delta h_{f,k}^0$ the standard formation enthalpies.

Crucially, from Eq. (2), it can be seen that to maintain a constant idealized reaction temperature level, $\frac{D\bar{T}}{Dt} = \frac{\partial\bar{T}}{\partial x_j} = 0$, subject to temperature sinks due to the endothermic reaction, the combination of pressure work and viscous dissipation must be designed to approximately balance the endothermic sink term $\dot{\omega}_T$ at each streamwise location, as illustrated in Fig. 7.

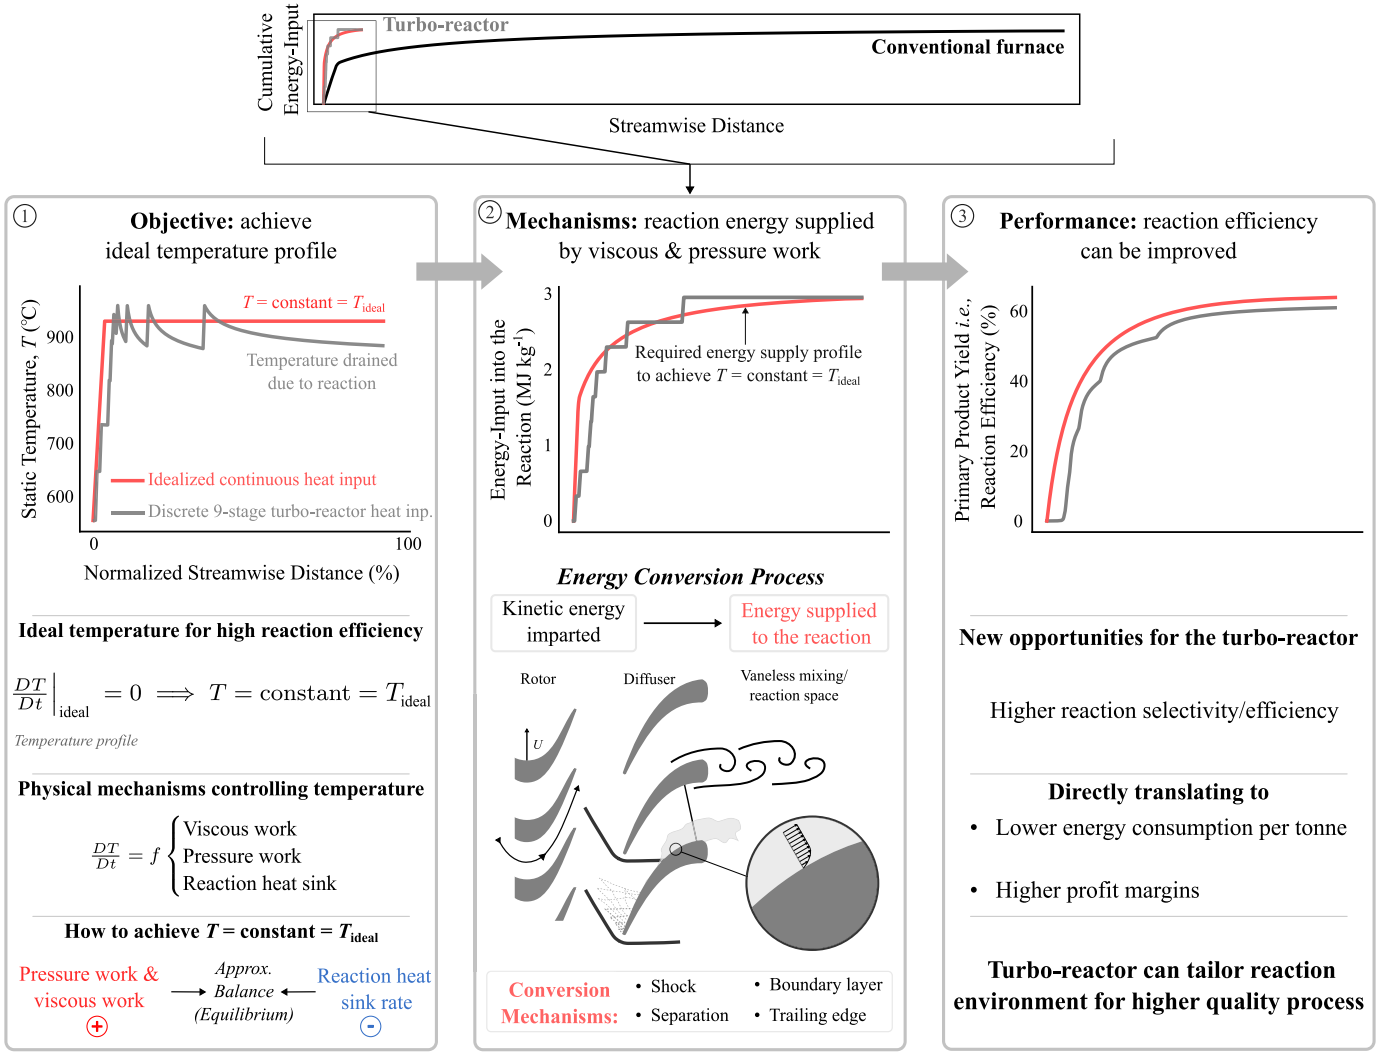


Fig. 7. Designing a reaction efficient temperature profile along the turbo-reactor bladed flow path.

3.1.2. Box ②: *Energy conversion mechanisms*

The aerodynamic designer is trying to shape the turbo-reactor flow path to achieve an energy supply profile to the reaction such that an ideal reaction-efficient temperature profile is matched. Box ② in Fig. 7 shows the energy conversion profile required to maintain some idealized temperature profile shown in Box ①. Through a series of discrete stages, the turbo-reactor attempts to approximate it.

There are reversible (isentropic) and irreversible (entropic) mechanisms responsible for this energy conversion profile. The first term in red in Eq. (2) is the reversible pressure work $\bar{u}_i \frac{\partial \bar{p}}{\partial x_i}$. This is a consequence of changes in the flow area and the reversible pressure rise component of the shock action.

The remaining energy conversion is provided by irreversible (entropic) energy conversion mechanisms through viscous shear work. These are accounted for by $(\bar{\tau}_{ij} + \tau^R_{ij}) \frac{\partial \bar{u}_i}{\partial x_j}$ in Eq. (2). Some of the primary physical mechanisms are highlighted at the bottom of Box ② in Fig. 7.

3.1.3. Box ③: *Reaction performance*

Controlling the balance between the endothermic sink and the dissipation profile at each streamwise location to better match an ideal reaction-efficient temperature profile can lead to improved reaction efficiency, as illustrated in box ③ of Fig. 7. To understand the importance of improving the reaction efficiency (*i.e.,* the yields), it is highlighted that for a typical hydrocarbon cracking plant (with capacity equal to 1 million tonnes of feed per year), a 10% improvement in reaction efficiency can increase annual profits by €60 – €90 million (as of 2024). This underscores the importance of tailoring the aerothermal design for improved yields. Note that this figure is calculated assuming a typical naphtha feedstock cost between 600 and 900 euros per tonne of feedstock [15].

It should be noted that the approach of striving for a *constant* ideal reaction temperature is not necessarily the current design strategy nor is it optimal for a general chemical reaction. However, it is a useful way to think about design goals: balancing viscous dissipation and reaction heat sink to produce a reaction-efficient temperature profile (see Fig. 7). For different feedstocks or reactions, the ideal temperature level may vary and/or may not be constant, leading to a nonequilibrium balance where $\frac{D\bar{T}}{Dt} \neq 0$.

3.2. *Theoretical limit on the rate of energy conversion*

To understand the potential range of temperature profile control, it is useful to determine the idealized energy conversion rate limit. This is set using a single normal shockwave with a high preshock Mach number (M_{preshock}), as shown in Fig. 8. The idealized energy conversion profile in Fig. 8 is determined from 1-D compressible flow equations with inlet properties determined from a meanline calculation. The supersonic flow is isentropically expanded (see Fig. 8) to a Mach number such that when it is terminated by a shock, the gas pressure is

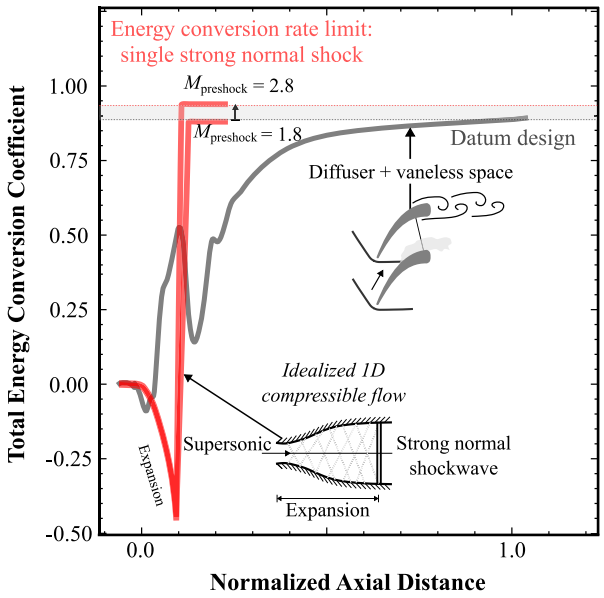


Fig. 8. A comparison of the total energy conversion coefficient (Eq. (6)) achieved with the datum design against the idealized limit transferable through a single normal shockwave within a 1-D nozzle.

raised to the fixed backpressure. It can be seen that over just a few millimeters, a single shock can dissipate over 90% of the incoming kinetic energy. However, the shock action yields diminishing returns with increasing M_{preshock}.

For rapidly preheating the gas mixture before the reaction phase, further capitalizing on shock action (see Fig. 8) could be a useful technique to reduce the flow path length by increasing the rate of energy conversion. However, in the reaction zone, which demands a specific temperature profile for maximum reaction efficiency, such a rapid rate of energy conversion may not be ideal for the reaction.

3.3. Defining energy conversion metrics

In conventional design practice, loss coefficients are used to quantify efficiency penalties and facilitate blade design comparisons. For the turbo-reactor design, the goal is to quantify the amount of incoming kinetic energy that is converted into enthalpy of the reaction. To quantify the irreversible (entropic) component of energy conversion, the relevant metric is the enthalpy loss coefficient used for conventional machines, as described by Denton [3]. However, in this case, we call it the "irreversible energy conversion coefficient". This coefficient is defined for a perfect gas as follows:

$$\zeta_{h,s} = \frac{c_p(T_{\text{exit}} - T_{\text{isentropic}})}{0.5V_1^2} \tag{4}$$

where subscript "1" defines the diffuser inlet and subscript "exit" defines the end of the vaneless space (see Fig. 9). The value of $T_{\text{isentropic}}$ is the static temperature obtained in an isentropic compression process that ends at the same final pressure as the actual process. By definition, the total energy conversion coefficient is defined as

$$\zeta_h = \zeta_{h,s} + \zeta_{h,\text{rev.}} = \frac{c_p(T_{\text{exit}} - T_1)}{0.5V_1^2}. \tag{5}$$

This represents the total change in enthalpy of the working fluid across the diffuser normalized by the incoming kinetic energy of the flow. By subtracting the irreversible energy conversion coefficient (Eq. (4)) from the total (Eq. (5)), the reversible (isentropic) component of the conversion coefficient must be defined as

$$\zeta_{h,\text{rev.}} = \frac{c_p(T_{\text{isentropic}} - T_1)}{0.5V_1^2} \tag{6}$$

4. Numerical methodology

A combination of unsteady Reynolds-averaged Navier–Stokes (URANS) and high-fidelity large-eddy simulation (LES) strategies has been used to illustrate the mechanisms responsible for the energy conversion process along the flow path. The following section (Section 4.1) describes the methodology developed to decompose the flowfield into the various mechanisms.

4.1. Energy conversion budget domain decomposition

Metrics

To decompose the energy conversion process into its constituent irreversible mechanisms, as shown in Fig. 9, a similar approach is taken to that of Pullan et al. [16] and Zhao et al. [17]. The domain is zonally decomposed into 11 spatial subvolumes intimately linked to the physical mechanisms discussed in Section 3.1. The Favre-averaged viscous dissipation rate (see Eq. (2)) is volume-integrated over each subvolume:

$$\Phi_\tau = \frac{\iiint_{\mathcal{V}} (\bar{\tau}_{ij} + \tau_{ij}^R)\frac{\partial \tilde{u}_i}{\partial x_j} d\mathcal{V}}{0.5\dot{m}V_1^2} \tag{7}$$

where $\mathcal{V}$ is the subvolume integrated over, $\dot{m}$ is the mass flow rate and V_1 is the diffuser inlet velocity. To quantify the reversible share of energy conversion, the pressure work term in Eq. (2) must also be volume-integrated:

$$\Phi_p = \frac{\iiint_{\mathcal{V}} \bar{u}_i \frac{\partial \bar{p}}{\partial x_i} d\mathcal{V}}{0.5\dot{m}V_1^2} \tag{8}$$

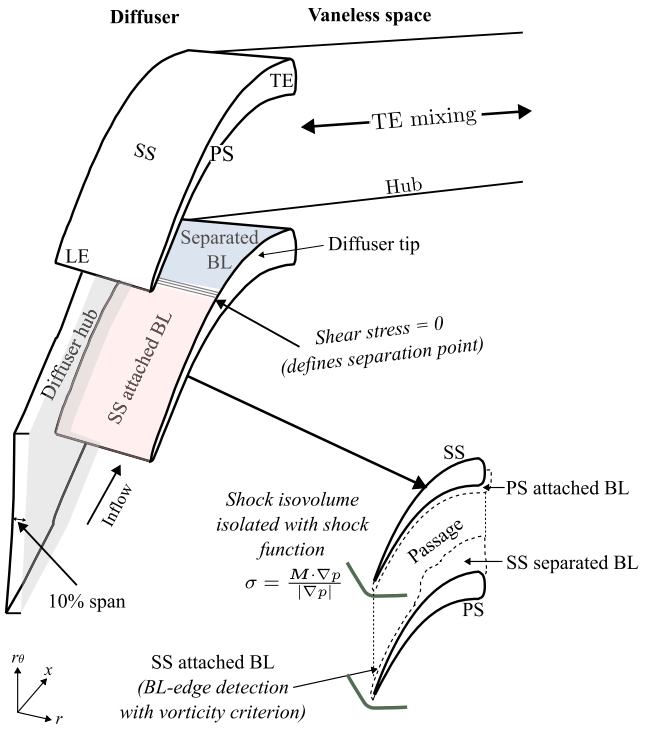


Fig. 9. A schematic illustrating the domain decomposition for the energy conversion breakdown analysis.

In conventional loss analysis, the term used to identify a source of loss is the entropy generation rate [3,18]. However, in the turbo-reactor, since the main objective of the energy conversion breakdown analysis is to identify sources that contribute to the conversion of kinetic energy into energy available for the reaction, the most relevant transport equation is the enthalpy (see Eq. (2)) rather than the entropy transport equation. Therefore, Eq. (2) implies that the viscous dissipation rate is the most appropriate metric.

Shock capturing

The isovolumes associated with a shockwave (see blue isovolumes at the Top of Fig. 12(*a*)) are identified using a criterion based on the shock function [19]:

$$\sigma = \frac{\boldsymbol{M} \cdot \nabla p}{|\nabla p|} \tag{9}$$

where $\boldsymbol{M}$ is the Mach number vector.

Boundary layer decomposition

The proceeding analysis distinguishes between an attached and a separated boundary layer. The boundary layer (BL) separation point is defined as the point where the wall shear stress falls to zero (see Fig. 9). The attached boundary layer edge at each streamwise location is determined as the pitchwise station with minimum radial vorticity.

Endwalls

The endwall flow at hub and tip is assumed to be all of the flow within 10% of the respective endwall. This includes the corner flow between the endwalls and the blade surfaces. Tip leakage is not included in this analysis.

Passage freestream mixing

The remaining mesh cells not yet classified are associated with passage freestream mixing.

Trailing edge mixing

In the vaneless space, all the non-endwall mesh cells immediately downstream of the diffuser trailing edge (TE) are associated with "TE mixing" (see Fig. 9).

4.2. Numerical setup

The in-house computational fluid dynamics solver TBLOCK [20,21] is used for all numerical investigations. The ability of TBLOCK to accurately predict the flowfield where the Reynolds number is similar to that observed in the turbo-reactor has been validated by Rosic et al. [22]. All numerical simulations are performed without a reaction model, and the working fluid is modeled as a calorically perfect gas. This assumption is reasonable since only a single stage is studied and thus the thermophysical fluid properties do not change significantly within this truncated domain. Furthermore, the density changes induced by the reaction are not large enough to have a noticeable effect on the energy conversion breakdown within a single stage. Therefore, neglecting aerochemical interactions is a suitable approach for this analysis.

LES: Full stage

The full-stage wall-modeled LES domain contains 226 × 152 × 152 nodes in the axial, spanwise and pitchwise directions within the blade passages (totaling 120 million grid points across the whole domain), leading to an average $y^+ = 4.1$ in the diffuser–vaneless-space system. The Smagorinsky subgrid-scale model with van Driest damping is employed. For fixed Mach number M and Reynolds number Re computations, it can be assumed that $Re = 700,000$ and $M = 1.2$, in which $Re = \frac{\rho_1 V_1 b_{\text{true}}}{\mu}$ based on the diffuser inlet conditions, $\rho_1 V_1$, and the true chord length is b_{true} and μ is the molecular viscosity.

Table 1
URANS mesh sensitivity for the entropy loss coefficient ζ_s and the irreversible $\zeta_{h,s}$ & total ζ_h energy conversion coefficients.

Blade passage grid sizes	ζ_s	$\zeta_{h,s}$	ζ_h
122 × 81 × 81	0.67	0.63	0.90
153 × 103 × 103*	0.65	0.63	0.89
183 × 123 × 123	0.65	0.63	0.89

URANS: Diffuser and vaneless space domain

For parametric studies on the influence of the Reynolds and Mach number on the energy conversion breakdown, a truncated domain with only diffuser and vaneless space has been used (see Fig. 9).

The inlet boundary condition is taken as the 2-D nonuniform time-averaged field from an unsteady full-stage calculation. From precursor numerical computations, it was found that periodic unsteadiness from the passing of the upstream rotor wakes has only a minimal impact on the energy conversion breakdown. This implies that the dissipation mechanisms respond linearly to wake-induced perturbations in diffuser inlet conditions. Therefore, the diffuser can be safely isolated for this analysis. This is useful because it allows the direct influence of Reynolds and Mach numbers on energy conversion to be decoupled from variations in the upstream flowfield entering the diffuser due to changes in the operating regime.

The following approach is used to vary the Mach number and Reynolds number independently. The Reynolds number can be set by varying the molecular viscosity to match the desired Reynolds number. The Mach number can be set by varying the isobaric heat capacity for a fixed ratio of specific heats.

Table 1 indicates sufficient mesh independence with a domain of 153 × 103 × 103 grid points in the axial, spanwise and pitchwise directions (totaling 18 million grid points), producing an average $y^+ = 5.4$ at the design point. The Spalart–Allmaras turbulence model is employed with wall functions. Due to a high turbulence intensity [4], fully turbulent boundary layers are used. Fig. 10 shows a comparison of a URANS and LES computation. The URANS calculation shows good agreement with high-fidelity LES, which provides confidence in the numerical results presented.

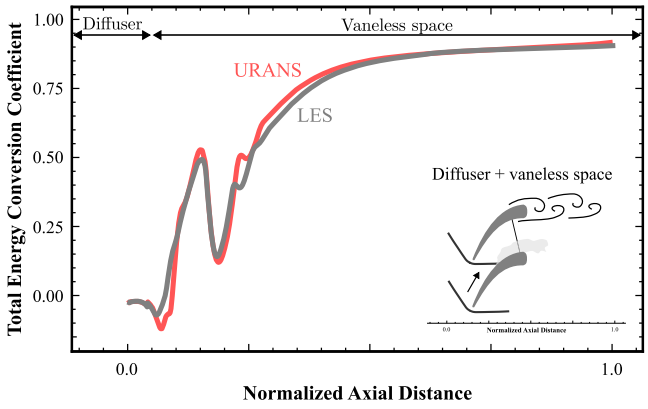


Fig. 10. A comparison of LES and URANS computations for the total energy conversion coefficient, ζ_h (see Eq. (5)) across the diffuser–vaneless-space system.

5. Energy conversion mechanisms

To enable fine-tuning of the energy conversion profile, it is important to understand the energy transformation processing mechanisms. This will inform future aerothermochemical design optimization to better match the optimal temperature profile for the chemistry (see Section 3.1). For example, the diffuser blade profile could be modified and/or the vaneless space length can be varied stage by stage (see Figs. 4 & 5).

5.1. Aerodynamically-tuned: Energy transfer process

Before the energy conversion process is illustrated in detail using numerical simulations, it is instructive to see how the aerodynamic losses (*i.e.,* the entropy rise) vary across the stator–rotor system (where mechanical energy is imparted). Fig. 11 highlights a clear distinction between the stator–rotor system and diffuser–vaneless-space system. While the latter is designed to amplify losses to chemically tune the reaction, the former is aerodynamically tuned, which is more in accordance with conventional turbomachinery design practice. However, while conventional design focuses on increasing efficiency across the rotor and throughflow capacity per unit frontal area of the machine, the design of the turbo-reactor stator–rotor system is more concerned with developing a high work coefficient as well as throughflow capacity. This directly translates into reducing flow deviation and blockage, which, of course, is strongly correlated with reducing the entropy rise (see Fig. 11). However, it is not necessarily the absolute minimum entropy rise that corresponds to the maximum power input.

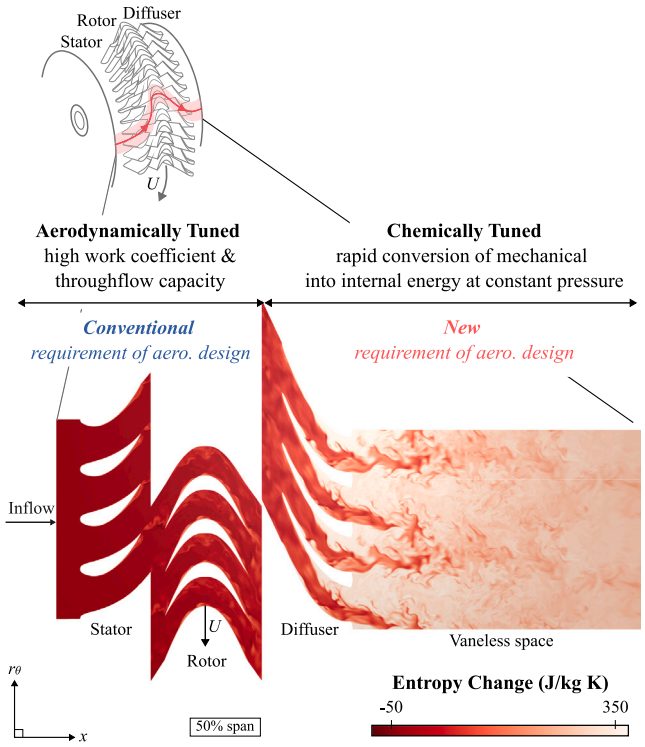


Fig. 11. Single stage instantaneous entropy change (LES).

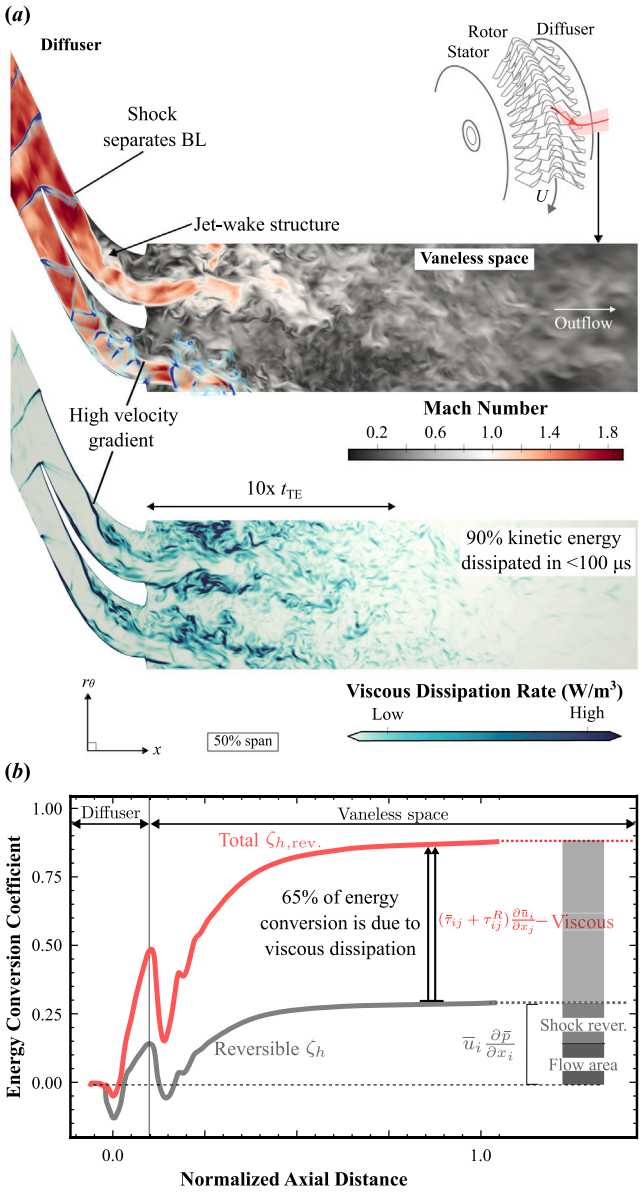


Fig. 12. (*a*) Instantaneous Mach number with superimposed numerical schlieren in blue (first) and viscous dissipation (second), and (*b*) time- and mass-averaged total (ζ_h) and reversible ($\zeta_{h,\mathrm{rev.}}$) energy conversion coefficient (LES).

5.2. Chemically-tuned: Energy conversion process

Fig. 12(*a*) shows the energy conversion process in the diffuser–vaneless system. Irreversibilities dominate over reversible pressure work, accounting for a 65% share (see Fig. 12(*b*)). This is primarily a consequence of turbulent rather than molecular viscosity. The diffuser blade row accounts for 55% of the total energy conversion. Within 10× diffuser TE thicknesses (t_{TE}), the energy conversion process is almost complete and, strikingly, ~90% ($\zeta_h \approx 0.9$) of the incoming kinetic energy is dissipated in under 100 μs (see Fig. 12(*b*)). The loss coefficient is approximately 10 times that developed in a compressor. Uniform conditions are quickly established just downstream of the diffuser TE, which is crucial for optimal reaction performance.

The remaining 35% of energy conversion is composed of fairly equal contributions from the reversible pressure gain component of the shock-induced temperature rise and isentropic diffusion of the flow (see Fig. 12(*b*)). Energy conversion across the thin shock layer has two distinct contributions (see Eq. (2)).

(1) Reversible pressure work (accounted for by $\overline{u}_i \frac{\partial \bar{p}}{\partial x_i}$).

(2) Irreversibilities (accounted for by $(\bar{\tau}_{ij}+\tau^R_{ij})\frac{\partial \bar{u}_i}{\partial x_j}$).

At low Mach numbers (M), the former dominates, but as M increases, the latter begins to dominate [23]. It will be shown in Section 5.3 that the temperature change due to entropy produced in the shock system is almost negligible. Therefore, the overall contribution of the shock system is about 20% (see Fig. 12(*b*)), primarily due to reversible pressure work rather than entropy rise.

Fig. 12(*a*) suggests that the dominant impact of the shock system on the irreversible energy conversion rate is an indirect rather than direct effect. The indirect impact of the shock system is the result of shockwave boundary layer interaction (SWBLI). Shock impingement separates the suction surface (SS) boundary layer at approximately 50% of the chord. The high-entropy boundary layer unravels and spreads, mixing out the hot boundary layer fluid into the freestream. At the

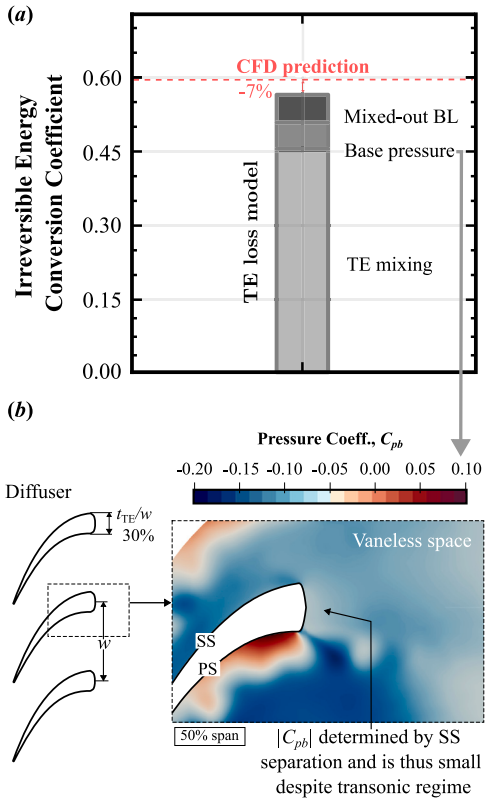


Fig. 13. (*a*) Irreversible energy conversion coefficient predicted by a modified Denton TE loss model [3,25] and (*b*) time-averaged pressure coefficient ($C_{pb} = \frac{p-p_{ref}}{p_{01}-p_{ref}}$, where p_{ref} is the average pressure PS and SS before the TE circle).

interface between the separated backflow region and the freestream jet above it, a large velocity gradient is generated (see Fig. 12(*a*)). This leads to intense deformation shear work ($(\bar{\tau}_{ij} + \tau^R_{ij})\frac{\partial \bar{u}_i}{\partial x_j}$) at the interface (see Fig. 12(*a*)). Furthermore, Fig. 12(*a*) (Top) shows, in blue, the footprint of a shock diamond within the overexpanded jet that exits the diffuser passage. This has important implications for TE mixing [24].

Despite the transonic flow regime and the thick TE occupying $t_{TE}/w \approx 30\%$ of the blade pitch w, surprisingly, Fig. 13 implies a minimal impact of the base pressure coefficient C_{pb} on TE loss. It is calculated as $C_{pb} \approx -0.05$ around the TE circle, which is several times smaller in magnitude than that of a transonic turbine ($C_{pb} \approx -0.3$) [24]. This is because, despite the global transonic flow regime, the flow locally around the TE circle is in the low subsonic regime and the pressure level is set by the SS separation (see Fig. 13(*b*)). Therefore, the magnitude of the base pressure is low, in the range experienced by subsonic blading.

Downstream, most of the mechanical energy dissipation is due to the mixing of two bulk streams — the jet and the separated wake (see Fig. 12(*a*)) — which have widely different stagnation pressures. The entropy generated in velocity equilibration to restore uniform flow is directly related to the difference in stagnation pressures between the jet and the wake [3]. Since this is large (see Fig. 12(*a*)), the downstream mixing intensifies. To model this mixing with a reduced-order model, a modified version of Denton's TE loss model [3,25] shown in Fig. 13(*a*) provides a fairly accurate estimate of the energy conversion despite the model being extrapolated substantially beyond its typical application range. Remarkably, the irreversible energy conversion coefficient is only underestimated by 7%.

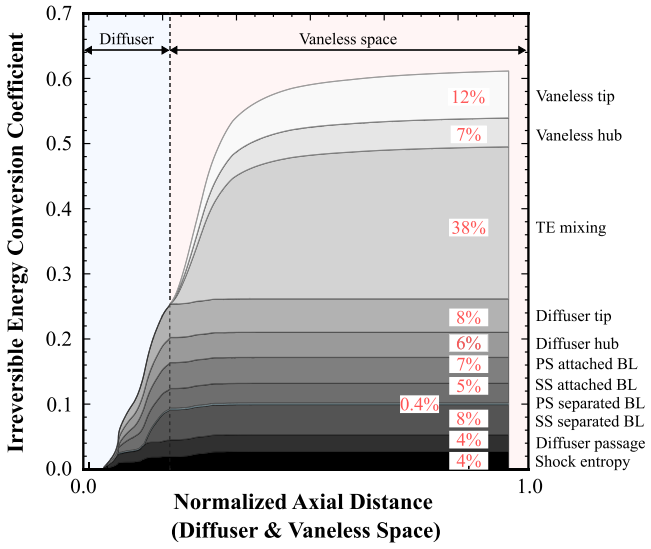


Fig. 14. Streamwise evolution of the irreversible energy conversion coefficient breakdown.

5.3. Global energy conversion breakdown analysis

Based on the domain decomposition introduced in Section 4.1, Fig. 14 shows the streamwise evolution of the irreversible energy conversion breakdown for a typical design and boundary conditions. Most importantly, it is clear that turbulent dissipation generated downstream of the diffuser TE is the dominant mechanism, accounting for 38% of the irreversible total. This process quickly reaches its maximum proportion within $1C_x$ (see Fig. 14). Because the SS boundary layer separates far upstream of the TE (see Fig. 12(*a*)), such a thick TE may not be necessary to encourage intense mixing (although it is beneficial from a thermomechanical perspective). However, in general, the TE thickness sets the turbulent length scale and the energy conversion rate scales with its thickness.

There are three further conclusions to be drawn from Fig. 14. First, the pressure-surface (PS) attached boundary layer energy conversion dominates over the SS BL because it remains attached longer. The SS boundary layer separates earlier (see Fig. 12(*a*)) and is therefore the most significant contribution within the diffuser passage. Second, in both the diffuser passage and vaneless space, the tip flow near the endwall generates more entropy than at the hub because the Mach number is higher and more mass flow is concentrated (which is implied by radial equilibrium). Finally, the viscous dissipation within the shock system is relatively small, as it mainly comprises relatively weak oblique shocks. This is consistent with the fact that the entropy change across a shock is only of third order in shock strength.

5.4. Mach and Reynolds number effects

For a perfect gas, three nondimensional flow parameters govern the energy conversion process in the diffuser: the Reynolds number, Mach number and diffuser inlet flow angle. This section explores Reynolds and Mach number effects over an engine-relevant range using four samples to cover each nondimensional group. This representative range covers the variability of nondimensional groups across the multistage environment [7]. Due to the unique incidence condition [26], the inlet flow angle is not constant as the diffuser inlet Mach number varies; however, it only varies by ~1°.

As the diffuser inlet Mach number increases for a fixed Reynolds number ($Re = 700{,}000$), Fig. 15 shows that the irreversible component of the energy conversion coefficient decreases while the reversible component increases. The reversible component increases with Mach

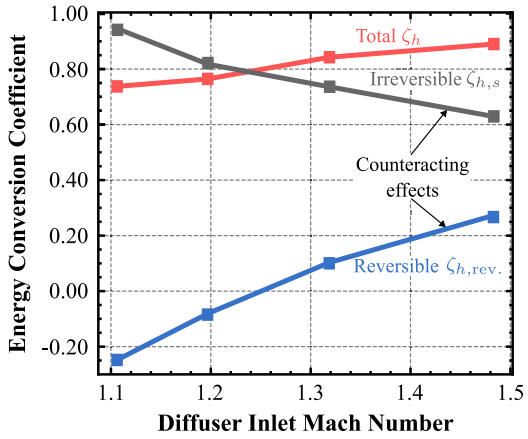


Fig. 15. Energy conversion coefficients as a function of diffuser inlet Mach number for a fixed Reynolds number.

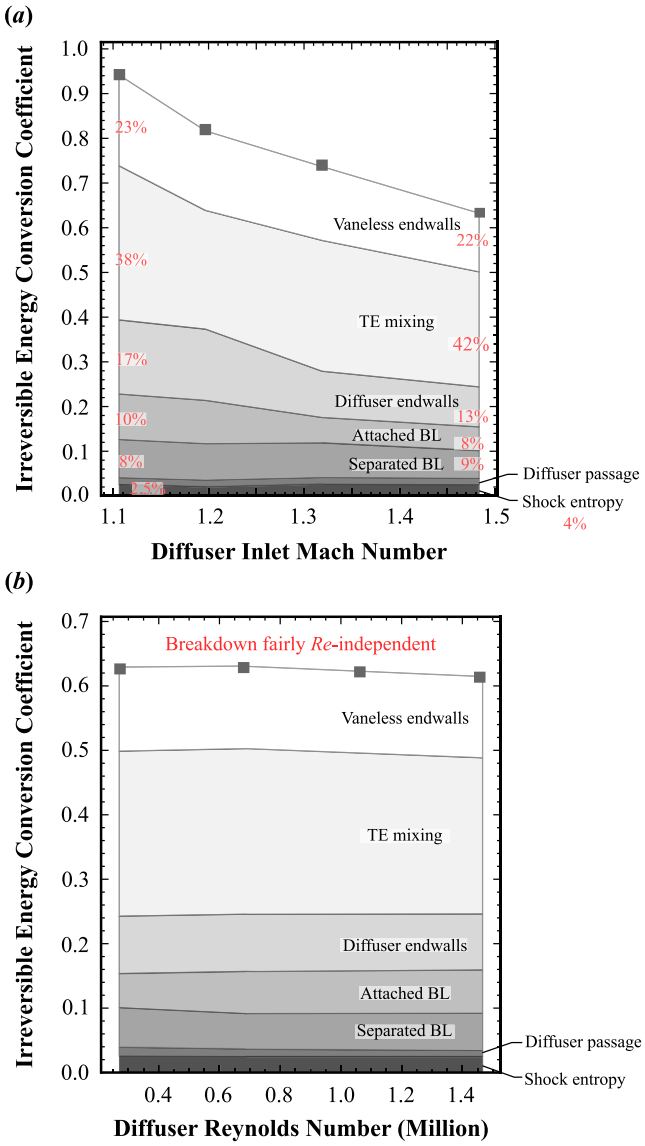


Fig. 16. The irreversible energy conversion coefficient as a function of: (*a*) diffuser inlet Mach number and (*b*) Reynolds number, $Re = \frac{\rho_1 V_1 b_{\mathrm{true}}}{\mu}$, based on the diffuser inlet conditions, $\rho_1 V_1$, and its true chord length, b_{true}.

number for two reasons. First, there is a greater contribution of the reversible pressure work through the shock system (see Fig. 15) due to an increased pressure ratio. Second, the passage isentropic flow diffusion increases due to changes in the effective flow area.

As a result of the counteracting dependence of the reversible and irreversible energy conversion coefficient with the Mach number, the change in the total energy conversion coefficient is weaker than expected with only a 20% increase in ζ_h within the range $1.1 \leq M \leq 1.5$. However, this effect should still be taken into account when designing a multistage machine in which the Mach number decreases towards the rear. This will influence the design of the reaction-efficient temperature profile.

Fig. 16(*a*) indicates that although the overall distribution of irreversible energy conversion mechanisms does not change dramatically, its overall magnitude decreases by 33% with increasing Mach number. A key factor is the reduction in normalized dissipation in the attached blade and endwall boundary layers. This is likely to be a consequence of two factors. First, the dissipation coefficient C_d is known to decrease with Mach number (for a fixed Reynolds number) [27], but only marginally. Second, it was found that at higher Mach numbers, the flow acceleration in the passage is reduced, changing the state of the boundary layer and leading to a lower C_d. Furthermore, the diffuser exit flow angle α was found to decrease by 13° over the range. This directly translates to lower TE mixing ($\zeta_{h,s} \propto 1/(w \cos \alpha)$), as seen in Fig. 16(*a*), as a result of reduced sudden flow expansion.

Compared to the Mach number, the consequence of Reynolds number variability is negligible and the energy conversion breakdown is approximately independent of the Reynolds number over a range from 0.4 to 1.4 million, as shown in Fig. 16(*b*). The Mach number is fixed at $M = 1.2$. There is a minor drop in the irreversible energy conversion coefficient at higher Reynolds numbers, since dissipation is inversely related to Re. Reynolds number independence is perhaps expected, since Re is above the critical Reynolds number ($Re_{\mathrm{crit}} \approx 100,000$) and turbulence intensity at the diffuser inlet is high (≥10%).

5.5. Compact multistage architecture

So far, the domain investigated has been a single stage with an extended vaneless space, which is required by the reaction. For preheating the nonreacting feedstock mixture before reaction activation, temperature profiling is unnecessary and a "compact" turbo-reactor stage can be adopted (see Fig. 4). Since the vaneless mixing space is absent, the flow can maintain a higher Mach number when it enters stages 2 to 4 relative to the first stage.

Fig. 17 illustrates that despite the removal of the vaneless mixing space by integrating stator and diffuser, it is still possible to dissipate a large amount of mechanical energy within each stage, albeit through different mechanisms. Whereas previously TE mixing dominated, now it no longer plays a significant role in the energy conversion process prior to the final diffuser. Instead, viscous dissipation within the blade surface and endwall boundary layers dominates. This is driven by an increasing freestream velocity across the multistage environment. Energy conversion near the endwalls is high as a result of additional strong secondary flows present here.

6. Summary and conclusions

This paper has explained that the turbo-reactor can be designed as a "chemically tuned" device in which the bladed path is tailored to optimize the reaction dynamics of a given feed. For the first time in

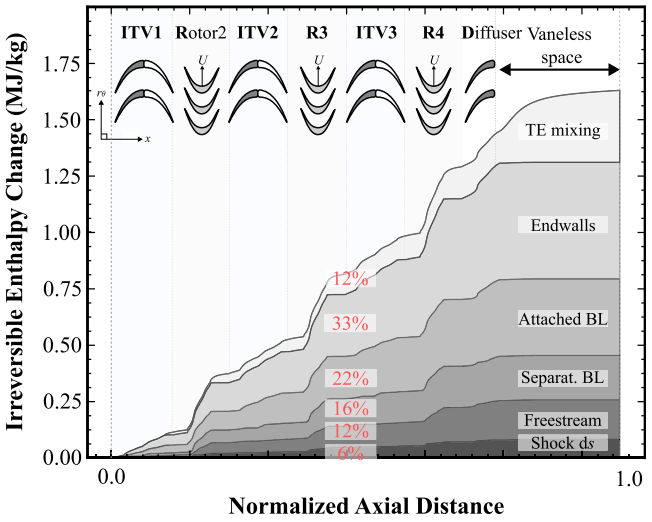


Fig. 17. Streamwise breakdown of the irreversible component of enthalpy change across the multistage compact turbo-reactor architecture.

turbomachinery, aerodynamic losses can be leveraged to control the rate of energy conversion into the reaction and thus design a reaction-efficient temperature profile. This can be achieved by optimally balancing dissipation against the endothermic sink term. By introducing new vectors of energy quality for turbomachines, that is, chemical exergy and reaction efficiency, the thermodynamic implications of converting electricity to heat have been clarified.

To enable optimization of the aerodynamic design to better match the ideal energy supply profile into the chemical reaction, it is crucial to explore the mechanisms responsible for the energy conversion process. A series of numerical investigations have been performed to examine the breakdown of energy conversion mechanisms. Several conclusions have been drawn.

(1) Irreversible energy conversion overshadows reversible pressure rise contributions, accounting for 65%. Of this 65%, turbulent viscous dissipation downstream of the diffuser trailing edge (TE) is the dominant entropic mechanism (at 25% of the total). This is also a critical process for rapidly homogenizing the reaction for improved performance.
(2) The total contribution of the shock system was around 20%, mainly due to a reversible pressure work mechanism rather than entropy rise. The dominant irreversible contribution of the shock system is an indirect one through shockwave boundary layer interaction. This process separates the boundary layer early in the passage and initiates a strong mixing process that continues downstream of the trailing edge.
(3) The energy conversion breakdown has been explored over a range of Mach and Reynolds numbers. Over a wide range of flow regimes it is possible to dissipate a large fraction of the incoming kinetic energy (≥60%). The energy conversion coefficient was independent of Reynolds number within a realistic range. As expected, the Mach number had a more significant influence. However, due to counteracting effects from reversible and irreversible energy conversion, the effect was relatively weak with only a 20% increase in the total energy conversion coefficient with Mach number rising from 1.1 to 1.5.
(4) For the compact stage design, the energy conversion rate remains high despite eliminating the interstage vaneless mixing space. This is partly a consequence of elevated skin friction due to increased flow speeds.

For the analysis presented, tip leakage was intentionally neglected. This is representative of a shrouded stationary blade row. However, a future study can incorporate leakages if necessary.

We hope that this paper can help to open an entirely new field of research into tuning the aerothermal design of turbomachines to control and optimize chemical transformations along the bladed flow path through careful temperature profile design.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Dylan Rubini was supported by funding from the Engineering and Physical Sciences Research Council (Grant number 2441788). The authors thank N. Karefyllidis for providing the stage dataset. The authors appreciate the fruitful discussions with N. A. Cumpsty. The computational resources used were provided by ARCHER2.

Appendix. Exergy

Total specific exergy is defined as follows:

$$b_{\text{ex}}^{tot} = b_{\text{ex}}^{Tp} + b_{\text{ex}}^{ch} \tag{10}$$

The physical (thermomechanical) component of the total exergy is:

$$b_{\text{ex}}^{Tp} = (h_0 - h_{0,\text{ref}}) - T_{\text{ref}}(s - s_{\text{ref}}) \tag{11}$$

where the subscript $\square_{\text{ref}}$ is the thermochemical state of the reference environment [28], h_0 is the stagnation enthalpy and s is the entropy. The chemical exergy is:

$$b_{\text{ex}}^{ch} = \sum_{k=1}^{N_{\text{spe.}}} Y_k b_{\text{ex},k}^{ch,ref} + RT_{\text{ref}} \sum_{k=1}^{N_{\text{spe.}}} X_k \ln X_k \tag{12}$$

where X_k are the species mole fraction, Y_k are the mass fractions, and $b_{\text{ex},k}^{ch,ref}$ is the chemical exergy of each species.

Chemical exergy is the result of the chemical disequilibrium of a species with the reference environment. This is calculated using the Gibbs function defined for a chemical reaction between the substance in question and the reference environment. In theory, useful work can be obtained from a chemical resource (species) by promoting a reaction until a chemical equilibrium is reached with the reference environment. The chemical exergy $b_{\text{ex},k}^{ch,ref}$ of various chemical species, fuels and materials can be found tabulated by Szargut et al. [28]. Calculating the chemical exergy of species is also discussed in Michalakakis et al. [29]. Note that for most hydrocarbon fuels, the chemical exergy is very similar to the heating value of the fuel.

Data availability

The data that has been used is confidential.

References

[1] J.H. Horlock, Advanced gas turbine cycles: a brief review of power generation thermodynamics, 2013.
[2] E. Michaelides, Exergy Analysis for Energy Conversion Systems, Cambridge University Press, 2021.
[3] J.D. Denton, The 1993 IGTI scholar lecture: Loss mechanisms in turbomachines, ASME J. Turbomach. 115 (4) (1993) 621–656, https://doi.org/10.1115/1.2929299, 10.
[4] D. Rubini, L. Xu, B. Rosic, H. Johannesdahl, A new turbomachine for clean and sustainable hydrocarbon cracking, ASME J. Eng. Gas Turbines Power 144 (2) (2021) 021024, https://doi.org/10.1115/1.4052784, 12.

[5] D. Rubini, N. Karefyllidis, L. Xu, B. Rosic, H. Johannesdahl, A new robust regenerative turbo-reactor concept for clean hydrocarbon cracking, J. Glob. Power Propuls. Soc. 6 (2022) 135–150, https://doi.org/10.33737/jgpps/150550.
[6] N. Karefyllidis, D. Rubini, B. Rosic, L. Xu, V.-M. Purola, A novel axial energy-imparting turbomachine for high-enthalpy gas heating: Robustness of the aerodynamic design, ASME J. Turbomach. 146 (3) (2023) 031005, https://doi.org/10.1115/1.4063928, 11.
[7] D. Rubini, N. Karefyllidis, B. Rosic, L. Xu, E. Nauha, Decarbonisation of high-temperature endothermic chemical reaction processes using a novel turbomachine: Robustness of the concept to feed variability, J. Glob. Power Propuls. Soc. 8 (2024) 111–126, http://dx.doi.org/10.33737/jgpps/185623.
[8] R. Chauhan, R. Sartape, N. Minocha, I. Goyal, M.R. Singh, Advancements in environmentally sustainable technologies for ethylene production, Energy Fuels 37 (17) (2023) 12589–12622, http://dx.doi.org/10.1021/acs.energyfuels.3c01777.
[9] O. Mynko, M. Bonheure, I. Amghizar, D.J. Brown, L. Chen, G.B. Marin, R. Freitas de Alvarenga, D. Civancik Uslu, J. Dewulf, K.M. Van Geem, Electrification of steam cracking as a pathway to reduce the impact of the petrochemical industry on climate change, J. Clean. Prod. 427 (2023) 139208, http://dx.doi.org/10.1016/j.jclepro.2023.139208.
[10] G.P. Thiel, A.K. Stark, To decarbonize industry, we must decarbonize heat, Joule 5 (3) (2021) 531–550, http://dx.doi.org/10.1016/j.joule.2020.12.007.
[11] M.E.H. Tijani, H. Zondag, Y. Van Delft, Review of electric cracking of hydrocarbons, ACS Sustain. Chem. Eng. 10 (49) (2022) 16070–16089, https://doi.org/10.1021/acssuschemeng.2c03427.
[12] T. Dickens, I. Day, The design of highly loaded axial compressors, ASME J. Turbomach. 133 (3) (2010) 031007, https://doi.org/10.1115/1.4001226, 11.
[13] M. Ghashghaee, R. Karimzadeh, Multivariable optimization of thermal cracking severity, Chem. Eng. Res. Des. 89 (7) (2011) 1067–1077, http://dx.doi.org/10.1016/j.cherd.2010.12.002.
[14] T. Poinsot, D. Veynante, Theoretical and Numerical Combustion, RT Edwards, Inc., 2005.
[15] IEA, Simplified Levelised Cost of Petrochemicals for Selected Feedstocks and Regions, IEA, Paris, France, 2017, https://www.iea.org/data-and-statistics/charts/simplified-levelised-cost-of-petrochemicals-for-selected-feedstocks-and-regions-2017.
[16] G. Pullan, J. Denton, E. Curtis, Improving the performance of a turbine with low aspect ratio stators by aft-loading, ASME J. Turbomach. 128 (3) (2004) 492–499, https://doi.org/10.1115/1.2182000, 03.
[17] Y. Zhao, R.D. Sandberg, Using a new entropy loss analysis to assess the accuracy of RANS predictions of an high-pressure turbine vane, ASME J. Turbomach. 142 (8) (2020) 081008, https://doi.org/10.1115/1.4046531, 07.
[18] M.B. Zlatinov, C. Sooi Tan, M. Montgomery, T. Islam, M. Harris, Turbine hub and shroud sealing flow loss mechanisms, ASME J. Turbomach. 134 (6) (2012) 061027, https://doi.org/10.1115/1.4006294, 09.
[19] D. Lovely, R. Haimes, Shock detection from computational fluid dynamics results, in: 14th Computational Fluid Dynamics Conference, 1999, p. 3285, https://doi.org/10.2514/6.1999-3285.
[20] J.D. Denton, An improved time-marching method for turbomachinery flow calculation, ASME J. Eng. Gas Turbines Power 105 (3) (1983) 514–521, https://doi.org/10.1115/1.3227444, 07.
[21] J.D. Denton, The use of a distributed body force to simulate viscous effects in 3D flow calculations, in: Volume 1: Turbomachinery, Turbo Expo: Power for Land, Sea, and Air, Vol. Volume 1: Turbomachinery, 1986, V001T01A058, http://dx.doi.org/10.1115/86-GT-144.
[22] B. Rosic, J.D. Denton, G. Pullan, The importance of shroud leakage modeling in multistage turbine flow calculations, ASME J. Turbomach. 128 (4) (2005) 699–707, https://doi.org/10.1115/1.2181999, 09.
[23] E.M. Greitzer, C.S. Tan, M.B. Graf, Internal Flow: Concepts and Applications, third ed., Cambridge University Press, Cambridge, 2007.
[24] J.D. Denton, L. Xu, The trailing edge loss of transonic turbine blades, ASME J. Turbomach. 112 (2) (1990) 277–285, https://doi.org/10.1115/1.2927648, 04.
[25] A.C. Senior, R.J. Miller, A data-centric approach to loss mechanisms, ASME J. Turbomach. 146 (4) (2023) 041007, https://doi.org/10.1115/1.4064167, 12.
[26] N.A. Cumpsty, Compressor Aerodynamics, Krieger Publishing, Florida, 2004.
[27] F. Tosto, A. Wheeler, M. Pini, Investigation of non-ideal effects in compressible boundary layers of dense vapors through direct numerical simulations, Phys. Fluids 36 (1) (2024) 016145, http://dx.doi.org/10.1063/5.0179570, 01.
[28] J. Szargut, D. Morris, F. Steward, Exergy Analysis of Thermal, Chemical, and Metallurgical Processes, Hemisphere Publishing Corporation, 1988.
[29] C. Michalakakis, J. Fouillou, R.C. Lupton, A. Gonzalez Hernandez, J.M. Cullen, Calculating the chemical exergy of materials, J. Ind. Ecol. 25 (2) (2021) 274–287, http://dx.doi.org/10.1111/jiec.13120.

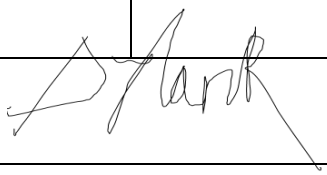
Statement of Authorship for joint/multi-authored papers for PGR thesis

To appear at the end of each thesis chapter submitted as an article/paper

The statement shall describe the candidate's and co-authors' independent research contributions in the thesis publications. For each publication there should exist a complete statement that is to be filled out and signed by the candidate and supervisor **(only required where there isn't already a statement of contribution within the paper itself).**

Title of Paper	Energy Conversion Mechanisms in Turbomachines for High-Temperature Endothermic Reactions: Redefining Energy Quality
Publication Status	□Published □ Accepted for Publication □Submitted for Publication □Unpublished and unsubmitted work written in a manuscript style
Publication Details	Dylan Rubini, , Budimir Rosic, and Liping Xu. "Energy conversion mechanisms in turbomachines for high-temperature endothermic reactions: Redefining energy quality".*Applied Thermal Engineering* 258 (2025): 124566.

Student Confirmation

Student Name:	Dylan Rubini		
Contribution to the Paper	Developed code for analysis and postprocessing and simulation, performed simulations, developed methodology, performed postprocessing and analysis and wrote manuscript		
Signature		Date	27/10/2024

Supervisor Confirmation

By signing the Statement of Authorship, you are certifying that the candidate made a substantial contribution to the publication, and that the description described above is accurate.

Supervisor name and title: Prof. Budimir Rosic		
Supervisor comments I confirm that the candidate's statement is an accurate description of his contribution to this publication.		
Signature 	Date	27/10/2024

This completed form should be included in the thesis, at the end of the relevant chapter.

3.1.1 Perspective

This chapter has reframed the turbo-reactor as a "chemically tuned" device, dispelling common misconceptions about the machine (*e.g.,* that it is a "bad compressor"). The RDR feeds a chemical transformation process by degrading the electrical exergy that drives the shaft. The objective is to maximise the amount of this exergy that appears as the chemical exergy of high-value species relative to undesired species. Since the electrical supply taps into a semi-limitless renewable energy source and has a considerably lower cost (see Fig. 1.5)) & exergy content than the feedstock, it is reasonable to sacrifice its energy quality in order to increase the efficiency of the chemistry. Since the cost of feedstock is high, a mere 1% improvement in reaction efficiency is equivalent to a $0.1 billion – $1.5 billion savings in annual feedstock costs for the industry.

As a result of the turbo-reactor's new set of requirements, there is a fundamental change in the design philosophy for turbomachines. Instead of curtailing aerodynamic losses, they are instead controllably elevated. Entropy generation can be exploited to design a reaction-efficient temperature profile by counterbalancing the enthalpy absorbed by the reaction with mechanical energy dissipation. For a typical turbo-reactor stage design, trailing edge (TE) dissipation—enhanced using a thick TE covering 30% of the pitch—is the dominant energy conversion source at 25%. This is followed by a 20% contribution from the shock system, primarily from pressure work rather than entropy production. In a real multi-stage machine, the Mach and Reynolds numbers decrease; it has been shown that the energy conversion profile is approximately Reynolds number independent but increases with Mach number because of the increasing strength of the shock system. However, the effect of Mach number was weaker than expected due to the counteracting effects of reversible and irreversible energy conversion mechanisms.

This chapter has introduced an entirely new paradigm within the aerothermal design space of turbomachines: the design of a reaction-efficient temperature profile. However, without a fast numerical design and analysis tool to capture interactions between aerodynamics and chemistry, we cannot feasibly optimise effectively within this design space. This once again highlights the importance of numerical tools for accelerated aerochemical coupling (which will be introduced in Chapter 6).

Chapter 4

Robustness of the Turbo-Reactor Aerodynamics

The underlying physics and fundamental design principles of the turbo-reactor guarantee that the flowfield is inherently more robust and controllable than other examples of energy-imparting machines, *i.e.,* compressors. As a consequence, the aerodynamic design can be flexible. This chapter explores two components of this flexibility.

4.1 Flexibility of the Axial Multi-Stage Arrangement

This section concentrates on the robustness of the aerodynamic concept to interstage mixing, and the final section (Sec. 4.2) focuses on the robustness of the concept to feed variability. In this section, high-fidelity numerical simulations will be used to determine the minimum interstage vaneless space length while still ensuring a well-conditioned and uniform flow entering successive stages. This unlocks the interstage vaneless space length as a tunable parameter to design a reaction-efficient temperature profile.

Pages 72–82:

A Novel Axial Energy-Imparting Turbomachine for High-Enthalpy Gas Heating: Robustness of the Aerodynamic Design

(*ASME. J. Turbomach.* 2023)

Check for updates

ASME Journal of Turbomachinery
Online journal at:
https://asmedigitalcollection.asme.org/turbomachinery

Nikolas Karefyllidis[1]
Oxford Thermofluids Institute,
Department of Engineering Science,
University of Oxford,
Oxford OX2 0ES, UK
e-mail: nikolas.karefyllidis@eng.ox.ac.uk

Dylan Rubini
Oxford Thermofluids Institute,
Department of Engineering Science,
University of Oxford,
Oxford OX2 0ES, UK
e-mail: dylan.rubini@eng.ox.ac.uk

Budimir Rosic
Oxford Thermofluids Institute,
Department of Engineering Science,
University of Oxford,
Oxford OX2 0ES, UK
e-mail: budimir.rosic@eng.ox.ac.uk

Liping Xu
Whittle Laboratory,
Department of Engineering,
University of Cambridge,
Cambridge CB3 0DY, UK
e-mail: lpx1@cam.ac.uk

Veli-Matti Purola
Coolbrook Oy,
Pieni Roobertinkatu 9,
Helsinki 00130, Finland
e-mail: velimatti@coolbrook.com

A Novel Axial Energy-Imparting Turbomachine for High-Enthalpy Gas Heating: Robustness of the Aerodynamic Design

Hard-to-abate industrial processes, such as petrochemicals and cement production, have long been considered technically challenging to decarbonize. In response to the urgent demand to eliminate industrial CO_2 emissions, a new class of energy-imparting turbomachines has been developed. These devices aim to convert mechanical into internal energy instead of pressurizing the gas, which enables high-temperature gas heating (up to 1700 °C) for a variety of applications. This article is organized into three parts. First, this article aims to demonstrate the capabilities of the novel, high-capacity, customizable, repeating-stage axial turbo-heater architecture for a hydrocarbon cracking example application. The study presents the new design requirements and working principles of this energy-imparting concept. The radically different objectives compared to a compressor enable ultra-high loading stage designs by avoiding the stability and efficiency constraints imposed on compressors. Within this new design space, the turbo-heater is able to achieve a loading coefficient $\psi \geq 4.0$. Second, detailed numerical simulations of a multistage axial turbo-reactor with various vaneless space lengths are conducted. This work conclusively demonstrates the robustness and flexibility of the aerodynamic design despite employing a uniform blade design. It is emphasized that the concept is tolerant to unsteady interstage turbulent disturbances, enabling nominal work input conditions to be achieved even for the most compact arrangements. Finally, having confirmed that the aerothermal restrictions on the vaneless space length can be removed, the designer is free to tailor the design to optimize the chemical reaction by (1) tailoring the residence time distribution to improve the yield and coking rate, (2) homogenizing reaction progress by mixing-out concentration gradients, and (3) adjusting the rotational speed to account for variations in the reaction dynamics for different feedstocks. [DOI: 10.1115/1.4063928]

Keywords: energy-imparting turbomachine, hydrocarbon cracking, high-temperature processes, decarbonisation, aerodynamics, energy transformation, computational fluid dynamics (CFD), turbomachinery blading design, large eddy simulation

1 Introduction

1.1 Motivation and Background. This article presents the development of a new class of energy-imparting turbomachines that addresses the engineering challenge of decarbonizing high-temperature heat within heavy industry. Historically, advances in the design of "energy-imparting" machines have been focused on improving cycle efficiency, widening operating range, reducing noise, mitigating emissions, etc., mainly for compressor systems in aeroengine and land-based power applications. However, the urgent need to eliminate industrial CO_2 emissions has led to a fundamental shift in objectives for these machines. Figures 1 and 2 illustrate that rather than converting the imparted mechanical energy into a pressure rise, the new concept, named the turbo-heater, aims to convert this work input into a static temperature rise (up to 1700 °C). To achieve this while maintaining an almost constant average static pressure p, maximizing the entropy rise $\left(\frac{T_2}{T_1} = \left[e^{\Delta s} + \left(\frac{p_2}{p_1}\right)^R\right]^{\frac{1}{c_p}}\right)$ is crucial.

Penetrating the high-temperature industrial market with a new turbomachinery-based technology capable of directly (and renewably) electrifying processes could contribute to eliminating more than 2 $BtCO_2$ emissions [1,2], while simultaneously lowering plant expenditure [3]. Projections indicate that with average renewable electricity prices falling below 30–40$/MW h and carbon prices rising above 100$/$tCO_2$, the low cost of electricity is set to trigger a rapid growth in industrial electrified capacity [4,5]. The new electric motor-driven turbo-heater could be the key enabler for decarbonizing challenging high-temperature sectors such as

[1]Corresponding author.

Contributed by the International Gas Turbine Institute (IGTI) of ASME for publication in the Journal of Turbomachinery. Manuscript received September 11, 2023; final manuscript received October 23, 2023; published online November 16, 2023. Tech. Editor: David G. Bogard.

 Copyright © 2023 by ASME

Downloaded from http://asmedigitalcollection.asme.org/turbomachinery/article-pdf/146/3/031005/7059051/turbo_146_3_031005.pdf by University of Oxford, Dylan Rubini on 30 January 2024

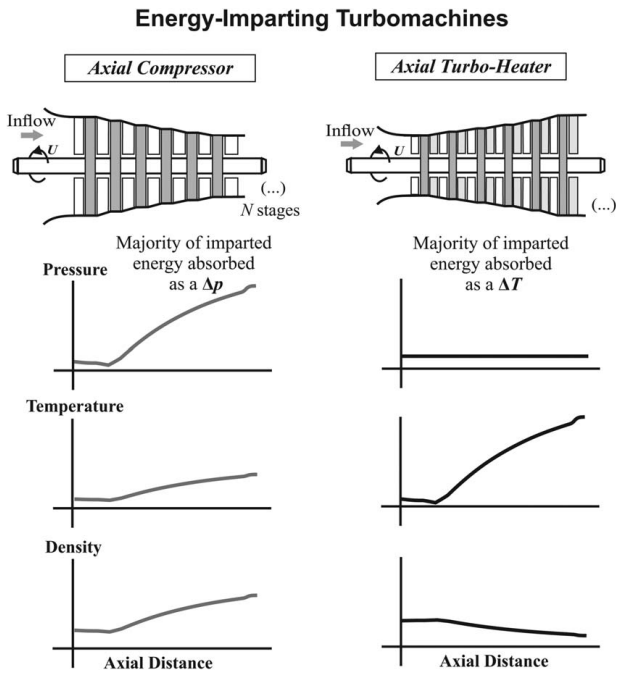


Fig. 1 Schematic comparison highlighting the key differences in the design and the thermodynamic profiles between an axial compressor and the new turbo-heater

steel and cement production, steam cracking, steam methane reforming, etc., in the "hard-to-abate" category.

Figure 3 illustrates that the turbo-heater technology can replace the inefficient and ineffective surface heat exchange energy transfer mechanism (in a furnace) with direct mechanical energy transfer into the working fluid (in a turbomachine) while dramatically increasing the power density (500 times). Compared to alternative pathways for decarbonizing hard-to-abate sectors, such as hydrogen-fueled and electric furnaces, an electric motor-driven turbomachine has three key advantages.

(1) Decades of experience with high-temperature and high-capacity gas turbine and compressor applications have provided accumulated know-how. In contrast, high-capacity and high-temperature electric heaters are still in their infancy.
(2) The conversion efficiency of hydrogen furnaces can be as low as 55% [6], whereas the turbo-heater solution boasts a conversion efficiency of up to 92%, as shown in Fig. 4.

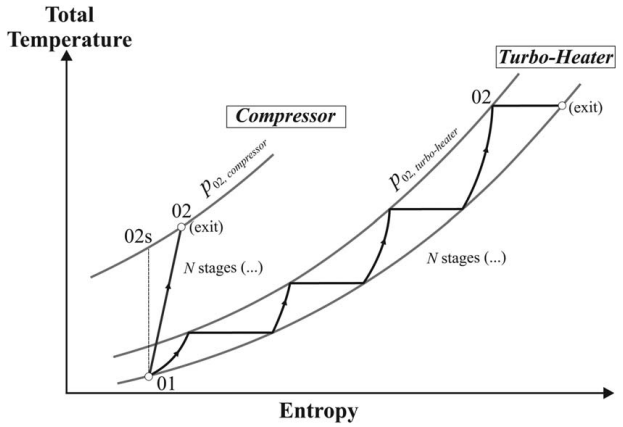


Fig. 2 An idealized schematic Mollier (or total temperature–entropy) diagram comparing a multistage axial compressor and the turbo-heater

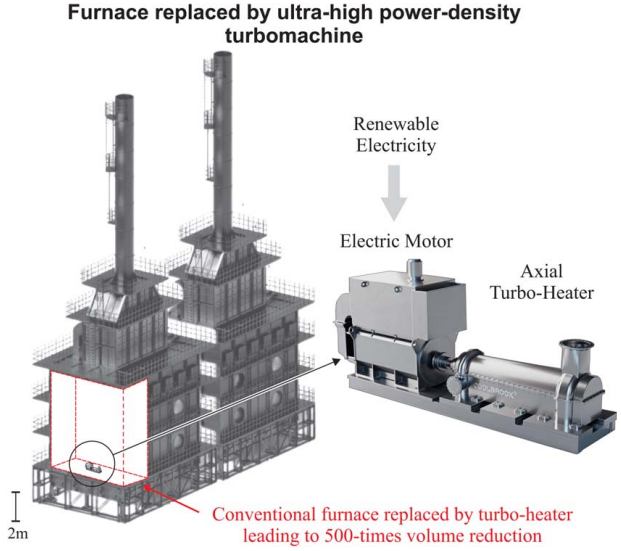


Fig. 3 A size comparison between an industrial cracking furnace and a scaled representation of the new axial turbo-heater [1] within the plant

(3) The fundamental change in the energy transfer mechanism of the turbo-reactor, when designed to drive a cracking reaction, could lead to substantial improvements in controllability, quality (e.g., higher yields and lower coking), and profitability relative to the conventional process. This is not possible with hydrogen or electric heaters.

The application of steam cracking for hydrocarbon feedstocks (such as ethane, naphtha, recycled plastic waste and bio-based feeds, etc.) at high temperatures is of significant interest due to the complexities of controlling chemical reactions along the gas path, which pose additional challenges for aerodynamic design. This article focuses on the development of a new turbo-heater specifically designed for this application, which is referred to as the turbo-reactor (or RotoDynamic Reactor (RDR) as shown in Figs. 3 and 5). The primary objective of the turbo-reactor is to drive chemical reactions, using a hydrocarbon–steam mixture as the working fluid to generate a series of endothermic reactions as part of a pyrolysis process. During this process, the long-chain, low-

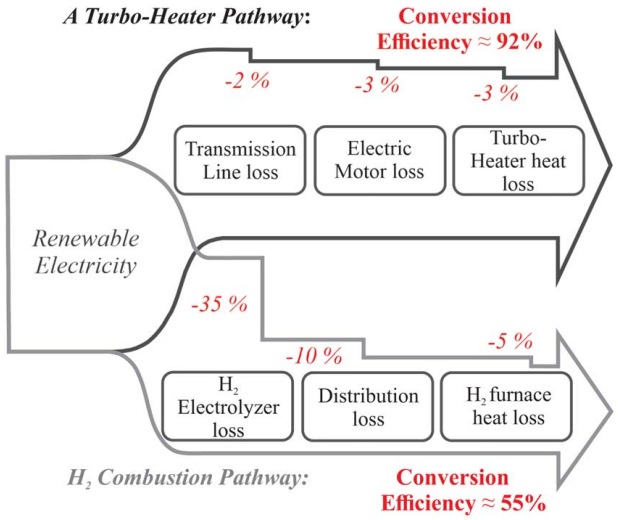


Fig. 4 Schematic comparison of the conversion efficiency for two potential pathways for decarbonizing high-temperature (900–1700 °C) industrial heat: direct electrification with the turbo-heater and hydrogen-fueled furnaces

Downloaded from http://asmedigitalcollection.asme.org/turbomachinery/article-pdf/146/3/031005/7059051/turbo_146_3_031005.pdf by University of Oxford, Dylan Rubini on 30 January 2024

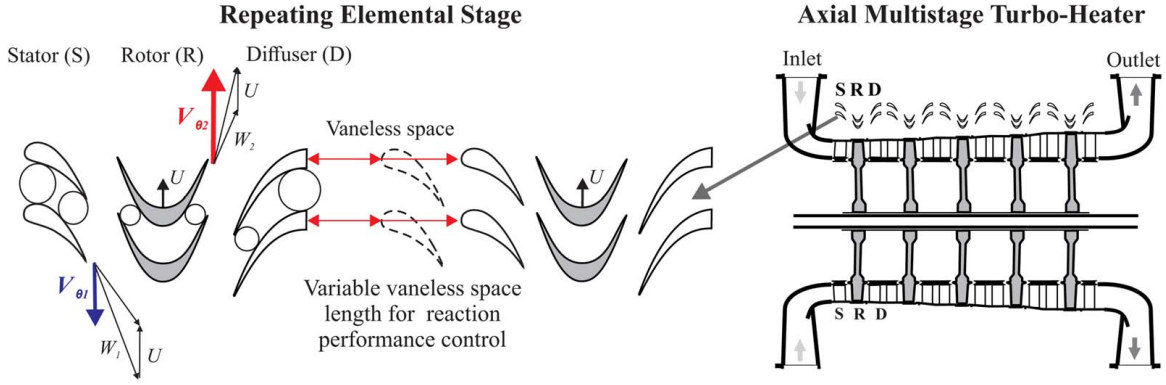


Fig. 5 A schematic of the axial turbo-heater including the velocity triangles

value feed undergoes thermal decomposition (cracking) into shorter-chain, higher-value products, such as ethylene and propylene. These basic chemicals are used to produce thousands of essential end-use products.

1.2 Scope of This Study. The primary objective and scope of this article is to first introduce a novel, customizable, repeating-stage axial multistage architecture for the RDR technology that is scalable and, therefore, well-suited for large-scale capacity. It is demonstrated that by removing the constraints imposed on axial compressors, a new ultra-high loading design space is opened. Second, by using high-fidelity numerical simulations, this work aims to conclusively demonstrate the robustness, flexibility, and controllability of this axial aerodynamic design. In spite of exploiting a simple, uniform, low-cost blade profile design, it is illustrated that the robust axial design is capable of tolerating unsteady interstage turbulent distortions. Even for the most compact multistage environment, the bulk rotor loading does not respond significantly to these interstage unsteady disturbances. This flexibility of the aerothermal design is the key feature of the axial concept that provides the designer freedom to optimize the residence time distribution and thus maximize the yield.

2 Unique Stage Design Objectives Enabling an Ultra-Power-Dense Turbomachine for Selective Control of Chemical Reactions

The novel design space boundaries surrounding the elemental stage design have led to a paradigm shift for ultra-high loading turbomachinery design (see Fig. 5). In light of a new application for turbomachines, the design objectives for conventional energy-imparting machines, specifically axial compressors, have been replaced. The objective of both compressors and turbo-heaters is to maximize the energy density imparted into the working fluid within a minimal number of stages. However, a significant difference exists between the two, while a compressor aims to minimize entropy generation and convert imparted energy into a pressure rise across a series of stages, the objective of the turbo-heater is to maximize the static temperature rise at a constant global static pressure (see Figs. 1 and 2). This shift in objectives relaxes the efficiency and stability constraints that limit the utilization of higher loading designs (where $\psi \geq 0.65$) from being used in axial compressors [7].

In the context of the turbo-heater, efficiency considerations primarily apply to the rotor, where maximum specific energy input is desired. Conversely, the impact of entropy generation in the stationary row on the performance and objectives of the turbo-heater is not significant. In comparison to axial compressors, a higher level of flow separation and entropy generation can be tolerated within the blade row downstream of the rotor. However, the level of flow separation that can be accommodated is limited by the requirement for a high throughflow capacity and, consequently, stage loading. Moreover, since the global pressure ratio remains constant, the adverse pressure gradient is eliminated, resulting in the removal of the onset condition for system instability. This factor further expands the ultra-high loading capabilities of turbomachines.

In the case of the axial gas heater used for steam cracking (i.e., the turbo-reactor), there are additional objectives necessary to optimize the pyrolysis reaction. For example, it is important to optimize the temperature profile for the reaction and eliminate nonuniformities.

The primary objective of constant-pressure ultra-fast gas heating concentrates the design space toward a radically new three-blade row repeating-stage design, as shown in Fig. 5. The stator is designed to accelerate and turn the flow, providing a supersonic, well-conditioned relative inflow to the rotor with a large tangential velocity component to maximize the work input potential. A supersonic impulse, ultra-high loading rotor blade row is used to increase the kinetic energy of the flow before immediately converting it into a static temperature rise by slowing down the high-speed flow, exploiting a shock train and promoting a mixing process within the diffuser blade row and interstage vaneless space. The interstage turbulent dissipation helps to maximize the fraction of imparted energy that is transformed into internal energy to raise the gas temperature rather than for pressurization. This is equivalent to maximizing the entropy rise, as illustrated in the Mollier diagram in Fig. 2. Within the interstage vaneless space and the highly accelerating stator, the flow resets and reconditions before the downstream stage. It is worth noting that for ultra-compact designs, the vaneless space shown in Fig. 5 may be entirely removed.

It is emphasized that through a fundamental switch in energy transfer mechanism (from heat exchange in a furnace to direct energy input in a turbomachine)—coupled with careful aerochemical tailoring of the elemental stage—it is possible to not only decarbonize a steam cracking unit but also to lower the cost, increase efficiency, and improve the chemical reaction performance. The energy transformation train enables three process quality and operability improvements relative to conventional furnaces:

(1) ultra-high enthalpy input rates that allow for a 500-fold increase in power density (Fig. 3)
(2) reduced coke formation rates due to substantially lower metal temperatures and fluid thermal and velocity gradients
(3) the ability to control temperature and residence time profiles to optimize the yields and reduce coking

This article focuses on feature (3), highlighting the principal advantage of the axial architecture for tailoring the residence time distribution to control the reaction. For the remainder of this article, the vaneless space may also be referred to the mixing space or reconditioning zone.

For further details on the elemental stage design and comparison with conventional furnaces, the reader is referred to the previous work by the authors [8–10] and the original patents [11,12].

Downloaded from http://asmedigitalcollection.asme.org/turbomachinery/article-pdf/146/3/031005/7059051/turbo_146_3_031005.pdf by University of Oxford, Dylan Rubini on 30 January 2024

3 Numerical Methods

The in-house computational fluid dynamics solver TBLOCK is used for all numerical investigations of the turbo-reactor concept [13–15]. The aerodynamic simulations are generally performed without a reaction model, unless explicitly stated otherwise (i.e., in Secs. 3.3 and 6), since their primary objective is to investigate the flow recovery between stages. The capabilities of TBLOCK at accurately predicting the flowfield, where the Reynolds number is similar to that observed in the RDR, were demonstrated in the work of Rosic et al. [16] for a multistage turbomachine.

A combination of high-fidelity large eddy simulations (LES) and scale-adaptive hybrid LES-URANS (unsteady Reynolds-averaged Navier–Stokes) calculations are employed. To aerodynamically investigate the elemental stage and interstage interactions, simulations of a two-stage LES domain (see Fig. 6) are conducted. For both the two-stage LES and the five-stage scale-adaptive simulation SAS, three different geometries were simulated by varying the vaneless space length (see Fig. 6). These domains include a 1, 2 and 4 $C_{x,\mathrm{DIF}}$ vaneless space arrangements. For a more detailed description of the datum geometry, numerical methodology, and operating conditions, the reader is referred to the prior work by the authors [8,9].

3.1 Two-Stage Large Eddy Simulations. The LES domain illustrated in Fig. 6 consists of 174 blocks with $128 \times 109 \times 109$ nodes in the axial, pitchwise, and spanwise directions within the blade passage and vaneless space (totaling to 205 million grid points across the whole domain).

The power spectrum of the streamwise velocity fluctuations shown in Fig. 7(*a*) indicates that the temporal resolution is sufficient to resolve the time scales of the higher-energy and lower-frequency eddies in the inertial subrange (up to $f \approx 2 \times 10^5$ Hz). Although a finer mesh and a higher temporal resolution could provide a more accurate representation of the highest frequency eddies and broaden the resolved spectrum, it is anticipated that this would neither affect the metrics of interest nor the conclusions drawn within this study. Since the flowfield is dominated by a strong blade pressure field, it is expected that the response of the flowfield within the blades is mostly overshadowed by the impact of large-length scale eddies. Figure 7(*b*) illustrates that the grid is capable of resolving a large fraction of the freestream turbulent kinetic energy, with over 90% resolved on average.

3.2 Five-Stage Scale-Adaptive Simulations. The scale-adaptive variant [17,18] of the Spalart–Allmaras (SA) turbulence model [19] is employed for all five-stage calculations. Figure 8 illustrates that the SA-SAS model accurately predicts the complex aerodynamics within the blade rows as well as the turbulent mixing decay rate in the vaneless space. This is the case despite using a ~8.5 times coarser mesh than LES calculations. The five-stage topology consists of 435 blocks with a mesh resolution of $92 \times 62 \times 62$ nodes in the axial, pitchwise, and spanwise directions within the blade passages.

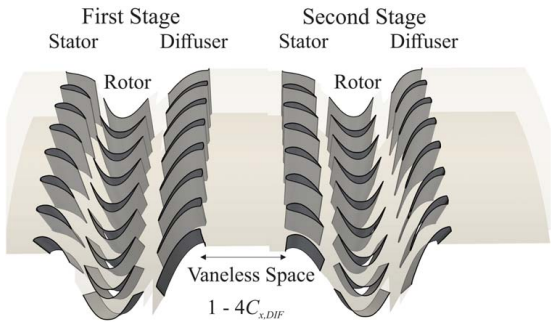


Fig. 6 Computational domain of a two-stage axial RDR

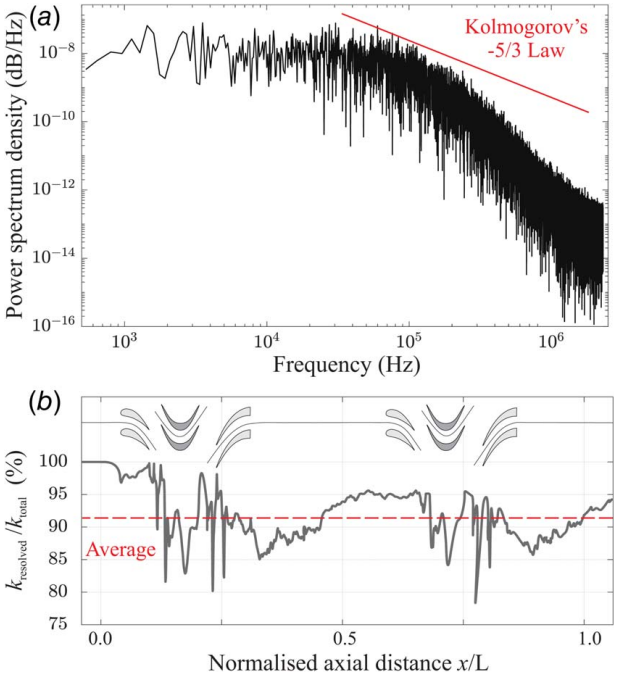


Fig. 7 (*a*) Turbulent kinetic energy cascade measured at 0.25% $C_{x,\mathrm{DIF}}$ downstream of the diffuser row and (*b*) percentage of resolved turbulence in the freestream

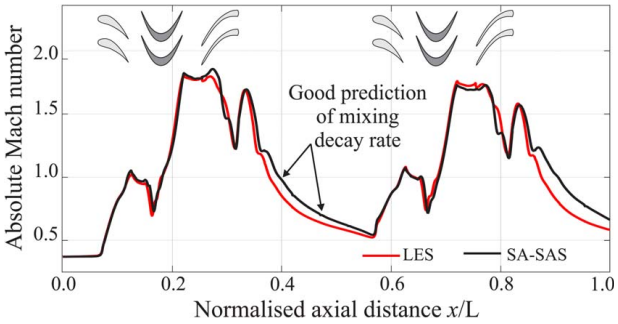


Fig. 8 A comparison of the time- and mass-averaged absolute Mach number distribution for the SA-SAS model against an LES solution for a two-stage reactor with 4 $C_{x,\mathrm{DIF}}$

3.3 Chemical Kinetic Modeling. Section 6 demonstrates how the reaction dynamics can be controlled through aerothermal design tailoring. Two levels of fidelity are employed for finite-rate chemical kinetic simulations, one with detailed reaction modeling in a 1D plug flow reactor (PFR) and the other with a reduced reaction mechanism (to maintain low computational cost) in a three-dimensional fluid-side modeling environment.

3.3.1 1D Plug Flow Reactor Model With Detailed Kinetics. Using RMG-PY [20], a detailed reaction mechanism is generated for a fixed naphtha feedstock composition, consisting of 220 species and 4000 reactions. An in-house 1D reactor model, built on top of a PFR model and stiff ODE solver within CANTERA [21], is used to approximate the enthalpy supplied by the rotor at fixed intervals within the convection time scale.

3.3.2 3D Reacting Flow Model With Low-Order Kinetics. In Sec. 6.1, a single-stage URANS calculation is used to illustrate

Downloaded from http://asmedigitalcollection.asme.org/turbomachinery/article-pdf/146/3/031005/7059051/turbo_146_3_031005.pdf by University of Oxford, Dylan Rubini on 30 January 2024

the mixing-out of species concentration gradients between stages. For this reacting flow simulation, a reduced chemical mechanism is employed since the primary focus is the mixing-out of the species with the highest mass fractions (e.g., ethylene). The composition at the inlet boundary comprises a mixture of naphtha and steam, with a steam-hydrocarbon ratio of 0.5.

4 Demonstrating the Working Principles of the Elemental Stage

This section demonstrates the working principles of the three-blade row stage design that enables a high level of specific energy to be imparted into the working fluid within a minimum number of stages.

4.1 A New Design Space for Ultra-High Rotor Loading. The stator blades, which function similar to nozzle guide vanes, are designed to accelerate, precondition, and pre-swirl the flow upstream of the rotor. Since two important constraints from compressor design are removed, it is possible to design for an ultra-high work coefficient. First, by removing the requirement for a pressure gain across the machine, entropy generation is less of issue than for compressors (other than in the rotor and stator). Second, the turbo-reactor has no adverse pressure gradient, and concerns about stall margin limiting the loading level are alleviated.

Ultra-high loading rotor blading is achieved using a supersonic (see Fig. 9) impulse, high-turning design. Figure 8 indicates that the flow enters with an absolute Mach number of 0.75 before being turned over 130 deg to discharge with a Mach number of 1.75. As a result, $\sim$300 kJ kg^{-1} of specific energy is imparted within a time scale of less than 0.1 ms. This is equivalent to a stage-loading coefficient of $\psi = 4.4$, an order of magnitude higher than a typical axial compressor. The impulse-type design has a stage reaction coefficient close to zero, corresponding to a negligible pressure rise across the row.

4.2 Rapid Energy Conversion Through Shockwaves. The diffuser row is multifunctional but primarily serves to de-swirl the flow exiting the rotor back into the axial direction and transform the imparted kinetic energy into internal energy. This process is initiated by a complex shock-train system, as shown in Fig. 9. By slowing down the high-speed flow and promoting interstage mixing (encouraged by shockwave boundary layer interaction), kinetic energy is converted into a static temperature rise within a fraction of a millisecond.

4.3 Importance of Interstage Mixing. Interstage turbulent mixing, as shown in Figs. 9 and 10, is essential to dissipate kinetic energy and homogenize the flowfield in the vaneless space. This turbulent mixing has two crucial design benefits: first, it ensures consistent velocity triangles for repeating-stage conditions, and second, it maximizes the entropy increase necessary for a constant-pressure rise in temperature. Approximately 60% of the imparted kinetic energy is converted into static temperature across the diffuser row, with an additional 25–30% dissipated in the interstage mixing space. For mixing spaces smaller than $1C_{x,\mathrm{DIF}}$, the flow would not be adequately reconditioned due to insufficient kinetic energy dissipation (<80%), leading to nonrepeating stage conditions.

4.4 Flow Recovery. To maintain a consistent specific energy input per stage, the highly unsteady, transonic, and turbulent flow exiting the diffuser must transition to a streamlined and settled state before entering the downstream stage. Any unsteady distortions that enter the stator increase the risk of inner passage flow separation and a subsequent drop in stage loading. Figure 10 indicates that the highest-energy mixing takes place within $0.5C_{x,\mathrm{DIF}}$ downstream of the diffuser before rapidly decaying to a turbulence intensity level similar to that of the clean reactor inflow. Due to the relatively low Reynolds number regime (50,000–200,000), disturbances are rapidly suppressed, ensuring that the flow has recovered to a sufficiently uniform state before entering the next stage.

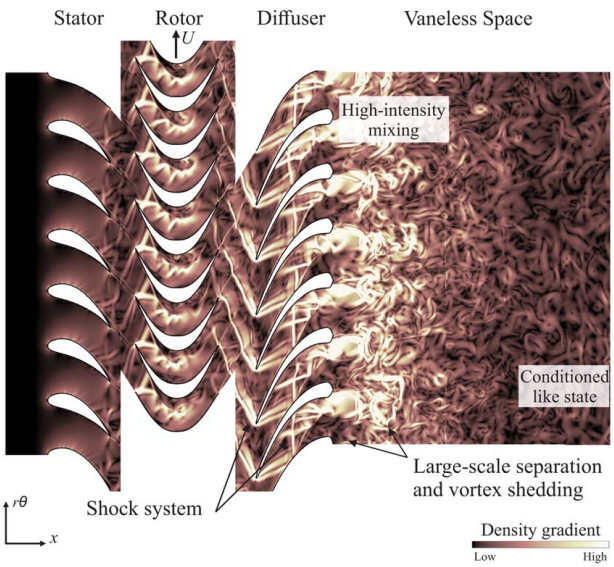


Fig. 9 Instantaneous numerical schlieren (density gradient) contours at midspan (LES)

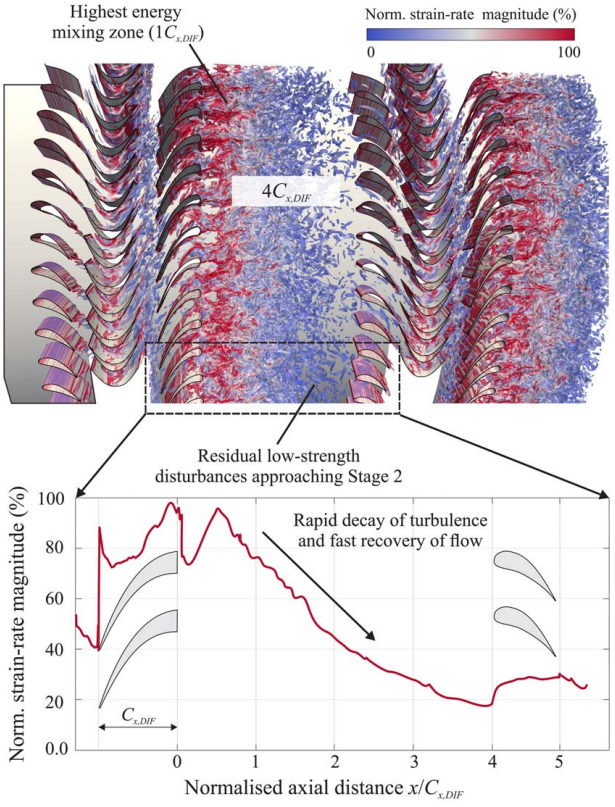


Fig. 10 Instantaneous Q-criterion isosurfaces across a two-stage axial RDR and time- and mass-averaged normalized strain rate magnitude line plot (LES). The strain rate is normalized referenced relative to inlet conditions.

Downloaded from http://asmedigitalcollection.asme.org/turbomachinery/article-pdf/146/3/031005/7059051/turbo_146_3_031005.pdf by University of Oxford, Dylan Rubini on 30 January 2024

5 Robustness of Multistage Aerodynamics

This section illustrates the robustness and flexibility of the multistage configuration with mixing space lengths ranging from 1 to 4 $C_{x,\mathrm{DIF}}$. The impact of interstage mixing is evaluated by comparing with a clean inflow scenario in the first stage, where no interstage mixing is present. Results show that a range of interstage flow-reconditioning-zone dimensions can be employed without compromising the performance. This allows for a tailored vaneless space length to be used for each stage, optimizing the residence time profile and enhancing the reaction yield.

5.1 Influence of Residual Disturbances on Downstream Stage Performance. Figure 11 shows that the ultra-compact arrangement ($1C_{x,\mathrm{DIF}}$) has the highest concentration of incoming residual vortical structures. As a result, Fig. 12 indicates variations in the spanwise relative yaw distribution. However, these spanwise differences in flow angle relative to the clean inflow case are reasonably small (within 5% at rotor inlet and outlet); therefore, nominal spanwise-averaged performance can be achieved even for the most compact arrangement.

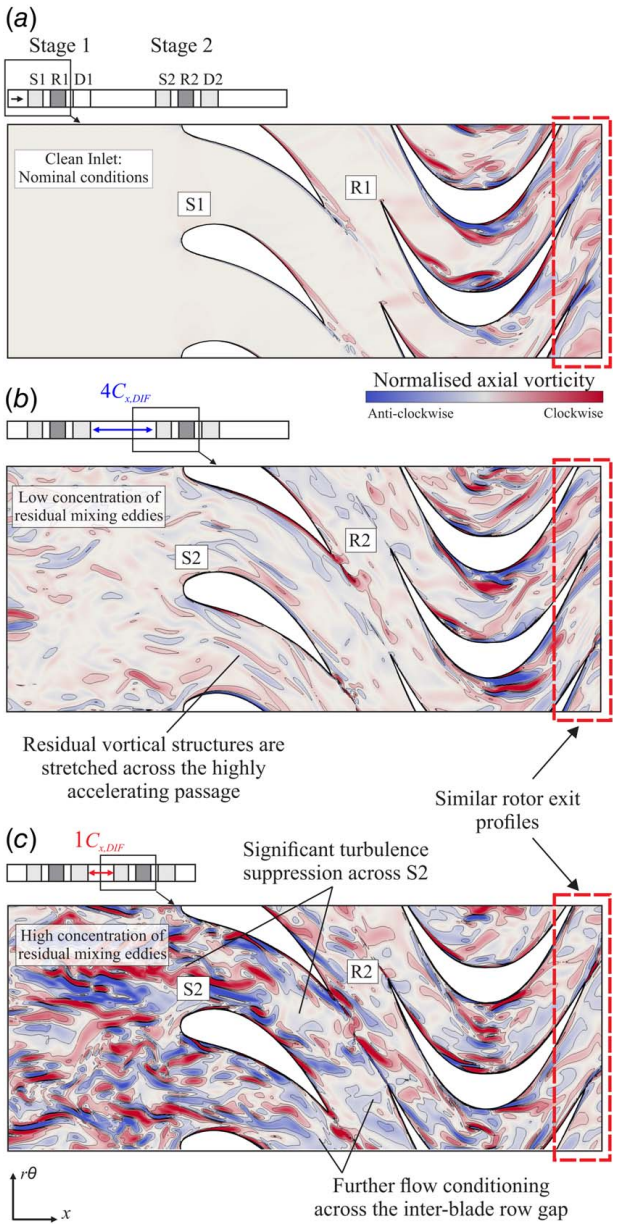


Fig. 11 Instantaneous axial vorticity contour at midspan for (*a*) clean inlet, (*b*) $4C_{x,\mathrm{DIF}}$ vaneless space, and (c) $1C_{x,\mathrm{DIF}}$ vaneless space (LES)

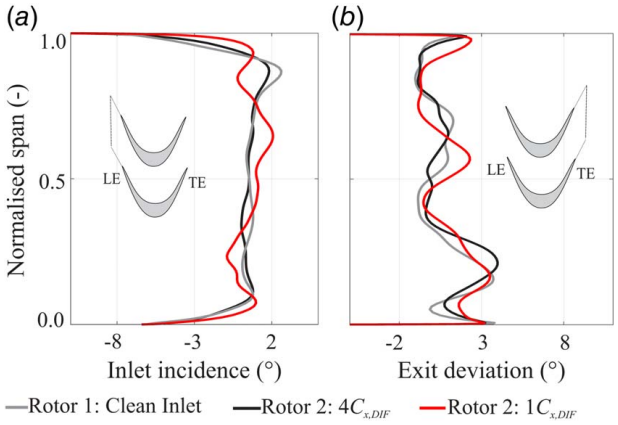


Fig. 12 Time- and pitchwise-averaged (*a*) rotor inlet incidence angle and (*b*) rotor exit flow deviation angle (LES)

Despite the varying levels of contamination entering the downstream stages, the similarity in rotor performance for different vaneless space lengths can be explained as follows.

(1) Figure 10 demonstrates that the disturbances have significantly decayed before entering the downstream stage due to a rapid rate of turbulent dissipation. This rate is a function of the low Reynolds number (50,000–200,000), Mach number regime, and diffuser trailing edge design, which influences the eddy size.
(2) The stator plays a crucial role in suppressing residual turbulence and damping unsteady incidence before the downstream stage, as shown in Fig. 11(*b*). Strong acceleration and flow confinement in the stator stretch and thin disturbances result in inlet incidence fluctuations being damped from ±7 deg to around ±2 deg for the $4C_{x,\mathrm{DIF}}$ test case and from ±20 deg to around ±5 deg for the $1C_{x,\mathrm{DIF}}$ test case at the exit of the stator. Despite the larger inlet incidence fluctuations for the $1C_{x,\mathrm{DIF}}$ test case, the stator is capable of adequately absorbing them, maintaining a similar range of flow deviation at the exit compared to the $4C_{x,\mathrm{DIF}}$ case.
(3) Comparing Fig. 13(*a*) with Figs. 13(*b*) and 13(*c*), interstage mixing does not significantly distort the time-mean blade loading distribution. The clean inflow test case in Fig. 13(*a*) illustrates fluctuations around the mean solely resulting from inner passage unsteadiness and blade row interactions. Figures 13(*b*) and 13(*c*) indicate that the superposition of additional fluctuations due to interstage mixing do not appreciably alter the blade loading relative to the clean case, as the disturbances have low amplitude and small length scales relative to the blade pitch. The rotor response is linear and quasi-steady, and the time-mean loading is not severely impacted. Thus, the change in enthalpy input between all the test cases is within 2%.

It has been demonstrated that despite an apparent increase in intensity of residual disturbances entering the downstream stage for the ultra-compact $1C_{x,\mathrm{DIF}}$ arrangement, the aerothermal profile is relatively unaffected compared with the clean inflow scenario. Therefore, ultra-compact vaneless space arrangements can be deployed for reaction optimization without penalizing the aerodynamic performance.

5.2 Single-Blade Profile Repeating-Stage Design. The feasibility of a low-cost, single-blade profile repeating-stage machine is investigated using five-stage numerical simulations. Figure 14(*a*) shows that each stage can consistently transfer kinetic energy into

Downloaded from http://asmedigitalcollection.asme.org/turbomachinery/article-pdf/146/3/031005/7059051/turbo_146_3_031005.pdf by University of Oxford, Dylan Rubini on 30 January 2024

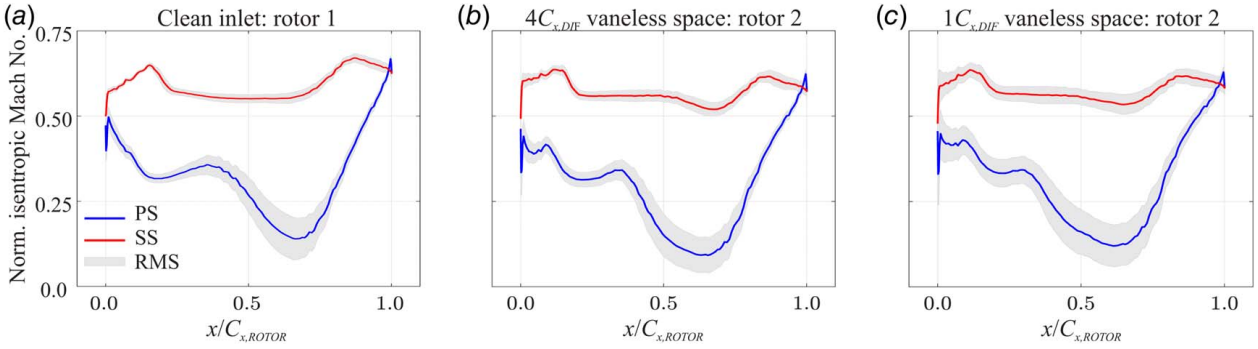


Fig. 13 Time-averaged normalized isentropic Mach number distribution (normalized by the exit Mach number at the rotor blade tip) at rotor midspan for (*a*) clean inflow (without interstage mixing), (*b*) $4C_{x,DIF}$ vaneless space, and (*c*) $1C_{x,DIF}$ vaneless space. It is noted that RMS is the root mean square of the fluctuations (LES).

internal energy, resulting in a nearly linear enthalpy rise across the machine at a constant overall pressure ratio (see Fig. 14(*b*)). Figure 15 demonstrates that with this uniform, repeating-stage design, it is possible to attain nominal specific enthalpy input regardless of the vaneless space dimensions.

The repeating-stage design performs well due to a nearly constant flow coefficient entering each of the five stages. This is possible because (1) the level of kinetic energy dissipation between stages is high and (2) temperature-driven density drop (40%) is relatively small compared to an compressor. Therefore, similar velocity triangles can be maintained across all stages without requiring expansion of the flow area. For scenarios with larger density variations, such as those with more stages or chemical reaction modeling, the meridional profile can be tailored accordingly.

5.3 Online Control of the Specific Energy Input. During online operation of the turbo-reactor, it may be necessary to adjust the specific energy level to recalibrate the cracking environment for different hydrocarbon feedstocks used in the plant, and maintain acceptable levels of feed conversion. The turbo-reactor provides wide operational flexibility, with relatively large adjustments to work input achievable through minor changes in rotational speed. Figure 16 illustrates the impact of a ±10% variation in rpm around the nominal design rotational speed Ω_{Design} on the total enthalpy rise across the machine, while keeping all other operational parameters, such as back pressure, constant.

The numerical results in Fig. 16 show that the enthalpy rise can vary by up to ±19% around the nominal conditions. This is in line with the ideal similarity rule, $\Delta h_0 \Omega^2$, which predicts a change in

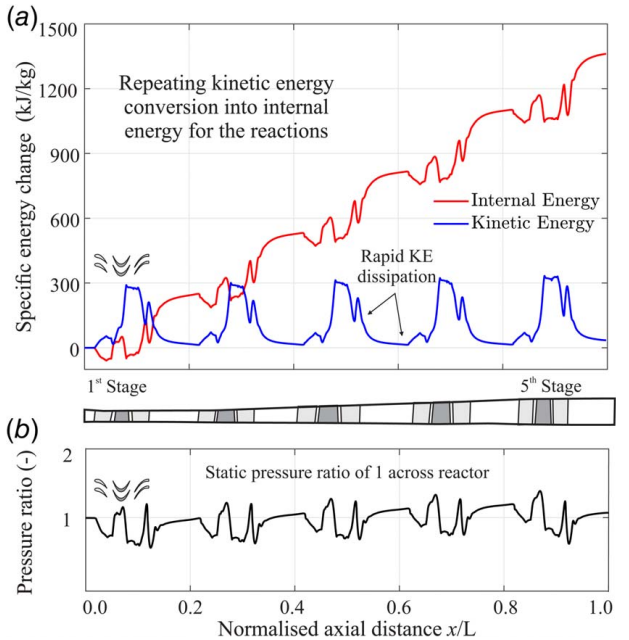


Fig. 14 Time- and mass-averaged line plot across five stages for (*a*) specific energy change relative to the inlet and (*b*) static pressure ratio (SA-SAS)

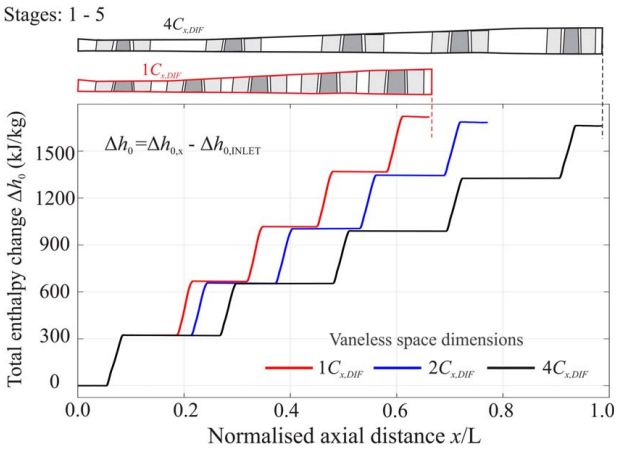


Fig. 15 Time- and mass-averaged total enthalpy change referenced to the inlet across a five-stage machine (SA-SAS)

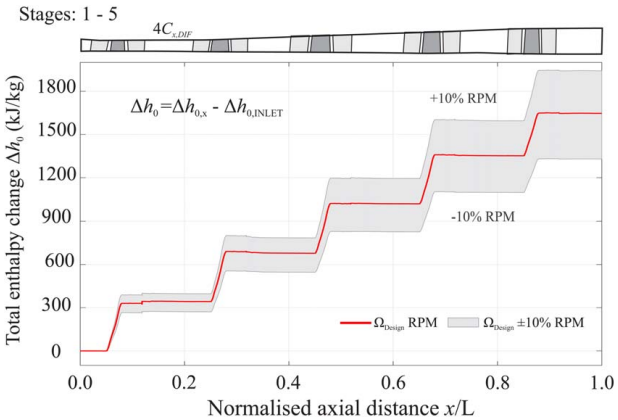


Fig. 16 Time- and massed-average operational envelope of the total enthalpy change relative to the inlet for Ω_{Design} ± 10% rpm around the design rotational speed (SA-SAS)

Downloaded from http://asmedigitalcollection.asme.org/turbomachinery/article-pdf/146/3/031005/7059051/turbo_146_3_031005.pdf by University of Oxford, Dylan Rubini on 30 January 2024

enthalpy of ±21%. However, the simulations also indicate a small deviation from the ideal rule due to minor distortions in the velocity triangles, leading to increased incidence, deviation, separation, and blockage. This is typical of a high-speed machine near the choke point where the mass flowrate no longer scales linearly with the blade speed. In this case, the capacity changes by ±4.5%, slightly less than the ±5% predicted by the ideal similarity rule.

6 Tailoring the Reaction Selectivity

The axial aerodynamic concept offers several opportunities to enhance reaction performance, i.e., to increase high-value product yields and decrease coke formation rate, compared to conventional radiant furnace technology. This section discusses these advantages.

6.1 Optimizing the Reaction Yield: Rapid Recovery of Homogeneity of the Species Concentrations. To maximize product yields and minimize coke formation rate, spatial uniformity of the velocity, temperature, and species concentration distributions are crucial in a high-yield and selective pyrolysis reaction, similar to removing hotspots in a combustor to minimize NOx emissions. Previous sections have shown that the aerothermal flowfield rapidly recovers between stages, and homogeneity is reset. However, it is equally important that the species concentration gradients are mixed out between stages, as shown in Fig. 17, to dampen any uneven progression through the reaction evolution in the spanwise and pitchwise directions. Spatial differences in reaction progress could amplify across a multistage machine because the chemical reaction is highly nonlinear. In the later stages, spatial regions that have progressed further will tend to undergo secondary side reactions that produce coke. Therefore, to avoid these situations and improve overall reaction performance, sufficient interstage mixing space is necessary, contributing to defining a lower bound on the vaneless space length.

6.2 Optimizing Reaction Yield: Residence Time Redistribution. The aerothermal properties of the turbo-reactor have been demonstrated to be robust and flexible, including for the ultra-compact vaneless space configuration ($1C_{x,\mathrm{DIF}}$), with rapid flow recovery between stages ensuring the work input potential is met for each stage. This relaxation of the aerodynamic barrier on minimum vaneless space length allows the designer to adjust the mixing space length and control the residence time distribution, reshaping the temperature profile for optimal reaction quality.

To determine the optimal residence time distribution subject to various physical constraints, a multiobjective genetic optimization algorithm is used within the open-source optimizer suite called PYMOO [22]. The design space comprises nine parameters defining

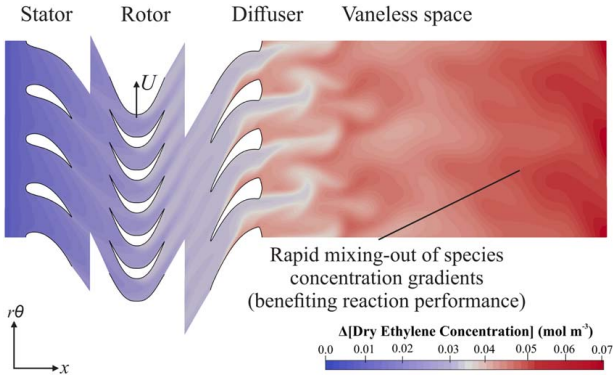


Fig. 17 Dry ethylene concentration contours at midspan indicating rapid damping of concentration gradients within a few diffuser chord lengths (reduced-order kinetics)

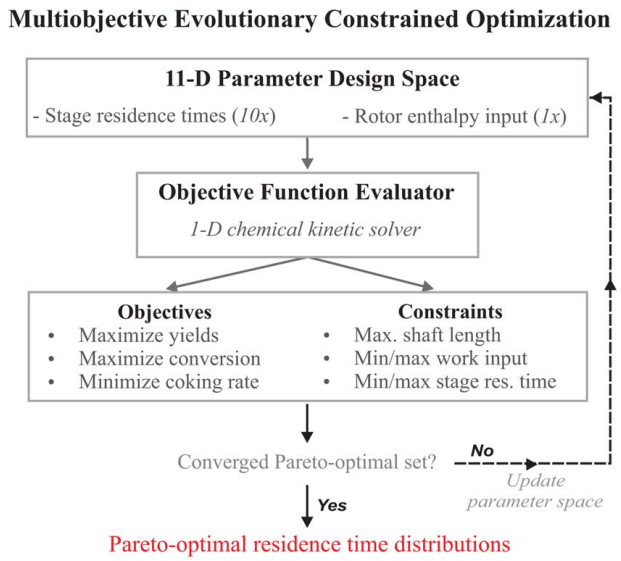


Fig. 18 A block diagram illustrating the constrained multiobjective optimization procedure for tailoring the residence time distribution

the mixing space length between each stage and one parameter defining the uniform rotor enthalpy input for all stages. To illustrate the advantages of redistributing residence time on the yield and coking rate, the number of stages is fixed at 9 during the optimization process. The objective functions and constraints applied during optimization are defined in Fig. 18, and the 1-D PFR model described in Sec. 3.3 is used to evaluate objective functions.

The performance of the optimized residence time distribution is evaluated by comparing it with the baseline design where the

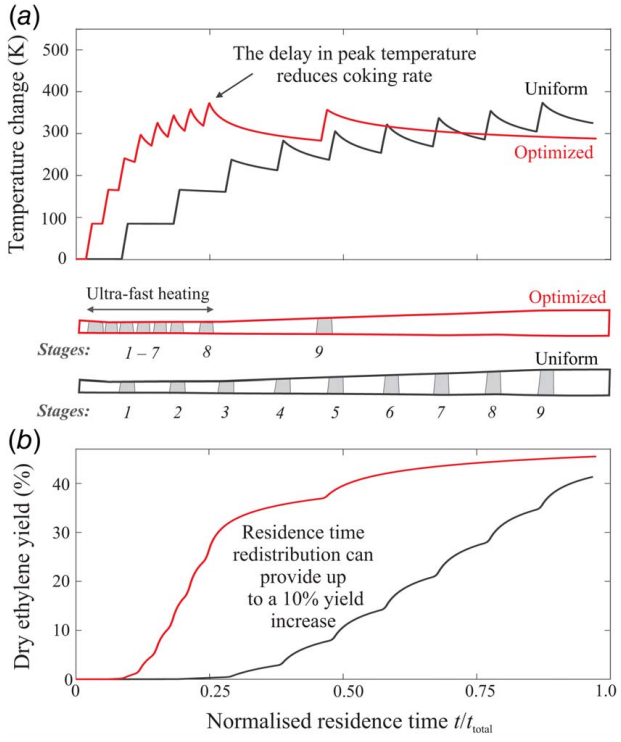


Fig. 19 A comparison of (*a*) the static temperature change relative to the inlet and (*b*) the dry ethylene yield for an optimized residence distribution reactor configuration against the uniform multistage design using a 1D PFR simulator

Downloaded from http://asmedigitalcollection.asme.org/turbomachinery/article-pdf/146/3/031005/7059051/turbo_146_3_031005.pdf by University of Oxford, Dylan Rubini on 30 January 2024

total residence time is distributed uniformly among the stages. The optimized and uniform cases use the same enthalpy input, number of stages, and overall residence time.

Figure 19 indicates that the optimized mixing space distribution favors residence times as low as aerodynamically possible for the majority of the forward stages before using "extended" residence time stages toward the rear of the reactor. This leads to a temperature profile where high-temperature conditions are reached very quickly and then held for the remainder of the reaction, which is inline with conventional wisdom [23]. The optimized profile enables a 10% higher primary product yield (see Fig. 19(*b*)) while simultaneously reducing coking rate by 10% (not shown for brevity). The maximum coking rate is reduced because the location of the peak temperature is shifted earlier and decoupled from regions of high ethylene concentration. Since coking is a strong function of both temperature and ethylene concentration, it is preferable to have lower temperature levels in regions where the ethylene concentration is high.

7 Summary and Conclusions

This study has introduced a novel, ultra-high power-dense, energy-imparting axial turbomachine that aims to decarbonize high-temperature, hard-to-abate industrial processes through direct renewable electrification. The concept's key features are demonstrated through high-fidelity numerical investigations, which reveal the following:

(1) Unlike an axial compressor, the primary objective of the turbo-heater is to raise the gas temperature (up to 1700 °C) at a constant pressure. As a result, the constraints on the stage-loading coefficient imposed by system instability and entropy generation are alleviated, enabling stable, ultra-high loading operation ($\psi \geq 4.0$). This ultra-fast energy input mechanism is complemented by an ultra-fast energy conversion process from kinetic energy into internal energy. This transformation of energy can take place over only a fraction of a millisecond by exploiting shockwaves and turbulent mixing.

(2) High-fidelity numerical investigations have suggested that the bulk, averaged blade loading and rotor performance do not respond to high-frequency, unsteady interstage turbulent eddies remnant from the upstream mixing process. The flow recovers and reconditions rapidly within only $1C_{x,\mathrm{DIF}}$ due to the relatively low Reynolds number regime, causing the length scales associated with unsteady disturbances to decrease rapidly and the amplitude to decay before the next stage. Second, the strong favorable pressure gradient in the stator ensures residual disturbances are strongly suppressed before the rotor. As a result, the fluctuations are damped to a sufficiently small amplitude that the rotor response can be approximated as linear and quasi-steady, and thus, the high-frequency fluctuations have little impact on the time-averaged loading. Consequently, the aerothermal performance has been shown to be independent of vaneless space length (within the tested range), which is confirmed for all stages of the multistage machine.

(3) Having conclusively demonstrated the flexibility and robustness of the axial architecture, this study has illustrated how the reaction dynamics can be optimized through multistage design tailoring and operational control without penalizing the aerodynamic performance. It is possible to simultaneously increase the primary product yield and minimize coke formation rates by (i) redistributing the residence time to tune the temperature profile, (ii) exploiting interstage mixing to promote a homogeneous reaction progress and therefore mitigate overcracking, and (iii) accounting for different feedstock characteristics by varying the rotational speed.

Acknowledgment

The authors would like to thank Coolbrook Oy for supporting this work and permitting the publication of this study. The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing facility in carrying out this work.

Conflict of Interest

There are no conflicts of interest.

Data Availability Statement

The authors attest that all data for this study are included in the paper.

Nomenclature

k = turbulent kinetic energy, $\mathrm{m^2 s^{-2}}$
$C_{x,\mathrm{DIF}}$ = diffuser axial chord length, m
$\psi = \dfrac{\Delta h_0}{U^2}$ = stage loading/work coefficient

Subscripts

x = axial component
$r\theta$ = circumferential component

Acronyms

LES = large-eddy simulation
RDR = RotoDynamic Reactor
RMS = root mean square
SA = Spalart–Allmaras
SAS = scale-adaptive simulation

References

[1] Coolbrook Oy, "Coolbrook Oy," Helsinki, Finland, https://coolbrook.com, Accessed October 15, 2022.
[2] Pörtner, H.-O., Roberts, D. C., Adams, H., Adler, C., Aldunce, P., Ali, E., and Begum, R. A., et al., 2022, "Climate Change 2022: Impacts, Adaptation and Vulnerability," *IPCC Sixth Assessment Report*.
[3] IRENA, 2022, "Smart Electrification With Renewables: Driving the Transformation of Energy Services," International Renewable Energy Agency, Abu Dhabi.
[4] Roelofsen, O., Somers, K., Speelman, E., and Witteveen, M., 2020, *Plugging In: What Electrification Can Do for Industry*, McKinsey, New York.
[5] Paltsev, S., Morris, J., Kheshgi, H., and Herzog, H., 2021, "Hard-to-Abate Sectors: The Role of Industrial Carbon Capture and Storage (CCS) in Emission Mitigation," Appl. Energy., **300**, p. 117322.
[6] Lu, Q., Zhang, B., Yang, S., and Peng, Z., 2022, "Life Cycle Assessment on Energy Efficiency of Hydrogen Fuel Cell Vehicle in China," Energy, **257**, p. 124731.
[7] Dickens, T., and Day, I., 2010, "The Design of Highly Loaded Axial Compressors," ASME J. Turbomach, **133**(3), p. 031007.
[8] Rubini, D., Xu, L., Rosic, B., and Johannesdahl, H., 2022, "A New Turbomachine for Clean and Sustainable Hydrocarbon Cracking," ASME. J. Eng. Gas Turbines Power, **144**(2), p. 021024.
[9] Rubini, D., Karefyllidis, N., Xu, L., Rosic, B., and Johannesdahl, H., 2022, "A New Robust Regenerative Turbo-Reactor Concept for Clean Hydrocarbon Cracking," Global Power Propulsion Soc., **6**, pp. 135–150.
[10] Rubini, D., Karefyllidis, N., Xu, L., Rosic, B., and Johannesdahl, H., 2022, "Accelerating the Development of a New Turbomachinery Concept in an Environment With Limited Resources and Experimental Data: Challenges," Presented at the ASME Turbo Expo 2022, Vol. 86120, Rotterdam, The Netherlands, June 13–17, American Society of Mechanical Engineers, p. V10DT36A001.
[11] Bushuev, V. A., 2016, "Bladed Reactor for the Pyrolysis of Hydrocarbons," Nov. 15, US Patent Application No. US 9.494,038 B2, US Patent and Trademark Office, WA, DC.
[12] Xu, L., and Rosic, B., 2020, "Rotary Device for Conducting Chemical Reactions," Aug. 18, US Patent Application No. US 10,744,480, US Patent and Trademark Office, WA, DC.
[13] Denton, J. D., 1983, "An Improved Time-Marching Method for Turbomachinery Flow Calculation," ASME. J. Eng. Power, **105**(3), pp. 514–521.

Downloaded from http://asmedigitalcollection.asme.org/turbomachinery/article-pdf/146/3/031005/7059051/turbo_146_3_031005.pdf by University of Oxford, Dylan Rubini on 30 January 2024

[14] Denton, J. D., 1986, "The Use of a Distributed Body Force to Simulate Viscous Effects in 3D Flow Calculations," International Gas Turbine Conference and Exhibit, Volume 1: Turbomachinery of Turbo Expo: Power for Land, Sea, and Air, Dusseldorf, Germany, June 8–12, p. V001T01A058.
[15] Denton, J. D., 1992, "The Calculation of Three-Dimensional Viscous Flow Through Multistage Turbomachines," ASME. J. Turbomach., **114**(1), pp. 18–26.
[16] Rosic, B., Denton, J. D., and Pullan, G., 2005, "The Importance of Shroud Leakage Modeling in Multistage Turbine Flow Calculations," ASME J. Turbomach., **128**(4), pp. 699–707.
[17] Menter, F., and Egorov, Y., 2005, "A Scale Adaptive Simulation Model Using Two-Equation Models," 43rd AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, Jan. 10–13, p. 1095.
[18] Coder, J. G., 2015, "A Scale-Adaptive Variant of the Spalart-Allmaras Eddy-Viscosity Model," 22nd AIAA Computational Fluid Dynamics Conference, Dallas, TX, June 22–26, p. 2462.
[19] Spalart, P., and Allmaras, S., 1992, "A One-Equation Turbulence Model for Aerodynamic Flows," 30th Aerospace Sciences Meeting and Exhibit, Reno, NV, Jan. 6–9, p. 439.
[20] Gao, C. W., Allen, J. W., Green, W. H., and West, R. H., 2016, "Reaction Mechanism Generator: Automatic Construction of Chemical Kinetic Mechanisms," Comput. Phys. Commun., **203**, pp. 212–225.
[21] Goodwin, D. G., Moffat, H. K., Schoegl, I., Speth, R. L., and Weber, B. W., 2022, "Cantera: An Object-Oriented Software Toolkit for Chemical Kinetics, Thermodynamics, and Transport Processes," https://www.cantera.org, Version 2.6.0.
[22] Blank, J., and Deb, K., 2020, "pymoo: Multi-objective Optimization in Python," IEEE Access, **8**, pp. 89497–89509.
[23] Van Goethem, M. W. M., 2010, "Next Generation Steam Cracking Reactor Concept," Ph.D. thesis, Delft University of Technology, Delft, the Netherlands.

Downloaded from http://asmedigitalcollection.asme.org/turbomachinery/article-pdf/146/3/031005/7059051/turbo_146_3_031005.pdf by University of Oxford, Dylan Rubini on 30 January 2024

Statement of Authorship for joint/multi-authored papers for PGR thesis

To appear at the end of each thesis chapter submitted as an article/paper

The statement shall describe the candidate's and co-authors' independent research contributions in the thesis publications. For each publication there should exist a complete statement that is to be filled out and signed by the candidate and supervisor **(only required where there isn't already a statement of contribution within the paper itself).**

Title of Paper	A Novel Axial Energy-Imparting Turbomachine for High-Enthalpy Gas Heating: Robustness of the Aerodynamic Design
Publication Status	☐Published ☐ Accepted for Publication ☐Submitted for Publication ☐Unpublished and unsubmitted work written in a manuscript style
Publication Details	Karefyllidis, N., Rubini, D., Rosic, B., Xu, L., and Purola, V. (November 16, 2023). "A Novel Axial Energy-Imparting Turbomachine for High-Enthalpy Gas Heating: Robustness of the Aerodynamic Design." ASME. *J. Turbomach*. March 2024; 146(3): 031005. https://doi.org/10.1115/1.4063928

Student Confirmation

Student Name:	Dylan Rubini		
Contribution to the Paper	Did optimization, methodology modelling, tool development and analysis for chemistry-related section of the paper. Write entire chemistry-related section of paper. Shared majority of analysis, postprocessing and writing workload for the entire paper.		
Signature		Date	16/08/2024

Supervisor Confirmation

By signing the Statement of Authorship, you are certifying that the candidate made a substantial contribution to the publication, and that the description described above is accurate.

Supervisor name and title: Prof. Budimir Rosic		
Supervisor comments I confirm that the candidate's statement is an accurate description of his contribution to this publication.		
Signature 	Date	19/08/2024

This completed form should be included in the thesis, at the end of the relevant chapter.

4.1.1 Chemically Tuning the Multi-Stage Arrangement

This section will further elaborate on the multi-stage arrangement optimisation and tuning procedure to achieve improved reaction efficiency and reduce the coke formation rate outlined toward the end of Sec. 4.1. The interstage separation was optimised for each stage to design a reaction-efficient temperature profile.

This section covers several aspects of chemically optimising the energy supply profile that were not addressed in the previous section. This includes varying the number of stages, relaxing the maximum peak temperature constraint, and allowing for a non-uniform enthalpy input level and rate for each stage.

Before detailing the multi-objective optimisation methodology, the 1-D chemical kinetic modelling methodology is introduced.

The Chemical Kinetic Mechanisms and 1-D Plug-Flow Reactor Model

This study concentrates on two hydrocarbon feedstocks: ethane and n-hexane. This covers a wide range of feed molecular weights, from 30 to 86 $\mathrm{g\,mol^{-1}}$, and hence reaction network complexity. A chemical kinetic mechanism (*i.e.,* a reaction network) is automatically generated for each feedstock using RMG-PY [87], containing 217 species and 3464 reactions for ethane and 113 species and 2869 reactions for n-hexane. The mechanisms produced are relatively simple and fast to evaluate, but there is uncertainty in their accuracy in predicting more trace species (*e.g.,* aromatics such as benzene, toluene, xylene, etc.) and thus providing more accurate coking estimates. The kinetic model can be easily refined using a comprehensive and detailed mechanism, such as that presented by Aerssens *et al.* [88], as a drop-in replacement. The mechanism in Ref. [88] contains molecules with a carbon number ranging from 1 to 7 and consists of 915 species & 9384 reactions. These mechanisms take longer to evaluate, and therefore their adoption is deferred to a later date.

The chemical kinetics are solved using an in-house `Python` code built on top of CANTERA [89] that has been developed during this work. A constant-area plug-flow reactor (PFR) model [90] is used as a 1-D approximation of the streamwise multi-stage energy conversion profile in the turbo-reactor. The PFR model assumes a homogeneous mixture at constant pressure. This is a reasonable assumption because the global pressure

ratio is approximately constant in the turbo-reactor (see Sec. 4.1). Figure 4.1 illustrates that each stage of the turbo-reactor is represented by a PFR with a "virtual" heat flux

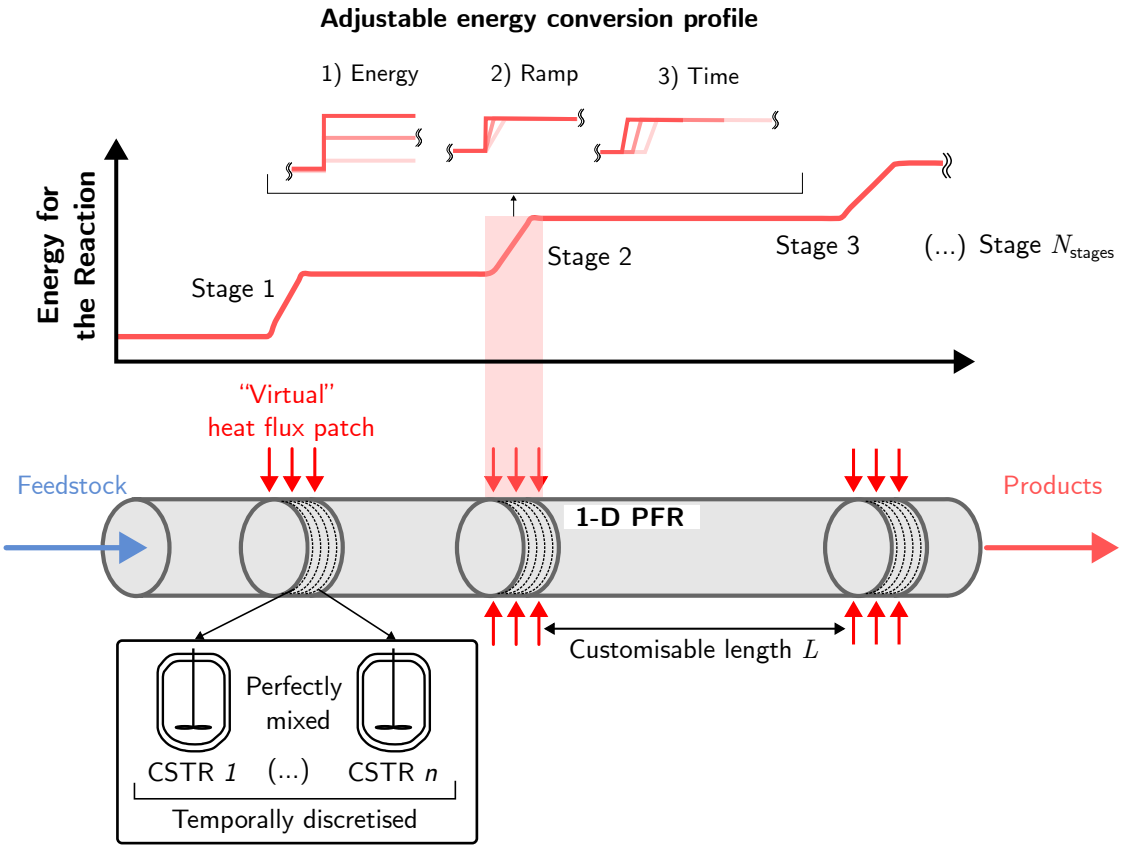


Figure 4.1: 1-D plug-flow reactor approximation of the multi-stage turbo-reactor used for chemical kinetic calculations

patch to approximate the energy supplied to the reaction. In the real machine, this would be due to the conversion of kinetic energy into internal energy.

The PFR model assumes that the species are perfectly mixed (*i.e.,* diffusional effects are ignored). Consequently, aerodynamics and chemistry are decoupled. The available control parameters are the heat flux magnitude, the heat flux ramp rate, and the residence time associated with each stage (see Fig. 4.1). Note that the heat flux ramp rate represents the rate of energy conversion from kinetic energy into internal energy.

Multi-Objective Optimisation Method

The goal of optimisation is to maximise the summed primary product yields (which is a combination of valuable olefinic species such as ethylene, propylene, 1-butene, 1,3-

butadiene, etc.) and minimise the coke formation rate subject to various constraints. These objectives conflict because, in general, higher temperatures are necessary to enhance yield, but this also leads to an exponential increase in the rate of coke formation. In this work, the problem is framed in a multi-objective rather than a single-objective environment, providing a non-dominated Pareto-optimal *set* of solutions [91]. The selection of a final solution from this set is deferred to a later date, once external auxiliary information is available to guide decisions.

Ideally, the two objective functions would be collapsed into one by converting both yield and coking into the same units using the equation proposed by Ghashghaee & Karimzadeh [92] and shown below:

$$f_c = 0.00001\phi_c F_c \left(\sum Y_k P_k - P_f Y_{\text{feed}}\right) - C_c\omega_c - \beta_c \tag{4.1}$$

where f_c is the gross profit ($\$\ \text{yr}^{-1}$), ϕ_c is the plant availability ($\text{h}\ \text{yr}^{-1}$), F_c is the hydrocarbon flow rate ($\text{kg}\ \text{hr}^{-1}$), Y_k is the effluent species mass fractions, P_k is the value of the products ($\$\ \text{t}^{-1}$), P_f is the feed cost ($\$\ \text{t}^{-1}$), C_c is the decoking cost ($\$$), ω_c is the number of decoking times (yr^{-1}) and β_c is the capital expense resulting from service life consumption caused by coking/decoking cycles ($\$\ \text{yr}^{-1}$). Unfortunately, the cost of decoking C_c and the influence on service life β_c are difficult to quantify at this stage of the design. This is why the multi-objective setting is more appropriate.

Multi-objective constrained optimisation is performed with a gradient-free evolutionary algorithm U-NSGA-III [93] within the Pymoo toolbox [91]. A gradient-free algorithm is preferred to a gradient-based algorithm because the parameter space can be large (up to 50 parameters) so the gradients of the loss function are costly to calculate and the loss function landscape is known to be multi-modal.

The first objective function is the sum of primary product species weighted by their market price. The second objective function is the peak coking rate, which sets the plant run-length. To evaluate the objective function at each iteration, the chemical kinetic model described above is used. The calculation of the coke formation rate is outlined in Appendix A. The boundaries of the input parameter space and the constraints applied

Table 4.1: Optimisation parameter space and constraints

Stage parameter	**Fixed/variable**	**lower bound**	**upper bound**
Energy input	variable	$100\,\mathrm{kJ\,kg^{-1}}$	$400\,\mathrm{kJ\,kg^{-1}}$
Residence time†	variable	0.1 ms	5.0 ms
Steam ratio	fixed	0.5	0.5
Pressure	fixed	2.0 bar	2.0 bar
Constraint	**Type**	**Value**	
Peak temperature	max.	850 – 1030 °C	
Feed conversion	min.	70 – 95%‡	
$\sum$ Residence time	max.	52 ms	

† there is a separate independent parameter per stage for residence time
‡ value within the range depends on the feedstock (*e.g.,* ethane or n-hexane)

during optimisation are detailed in Table 4.1. Unless stated otherwise, the residence time is optimised for each stage independently, but the level of energy supplied to the reaction is optimised as a single parameter uniformly for all stages. A constraint is imposed on the minimum acceptable feed conversion, since this is an important customer requirement based on the feed recycling capacity of the plant. For lighter feed, ethane, the minimum feed conversion required is 70% and the unconverted feed is recycled back to the inlet [33]. For heavier feed, n-hexane, no recycling is used, so the minimum feed conversion required is 95% [33]. For simplicity, the operating pressure and steam ratio were fixed, as their effect is of second order. The non-dominated Pareto front convergence after ∼22,500 model evaluations. This takes approximately 3.7 hours on an 18-core CPU.

Pareto Fronts: Number of Stages

A basic decision during the preliminary design of the turbo-reactor is setting the number of stages to distribute energy input over. Therefore, this section evaluates how much the Pareto front improves as the number of stages increases. Intuitively, having more stages implies that there are more control points to better design the optimum temperature profile for the reaction.

Figure 4.2(a) shows that increasing the number of stages from 10 to 16 shifts the optimal Pareto front toward a region of increased yield and lower coking. Fewer stages were also tested ($N_{\mathrm{stages}} \in [4, 7]$); however, these could not satisfy the maximum

temperature ($T_{\max} \leq 1000\,°\text{C}$) and minimum conversion ($C_{\text{feed}} \geq 70\%$) constraints and are therefore not shown.

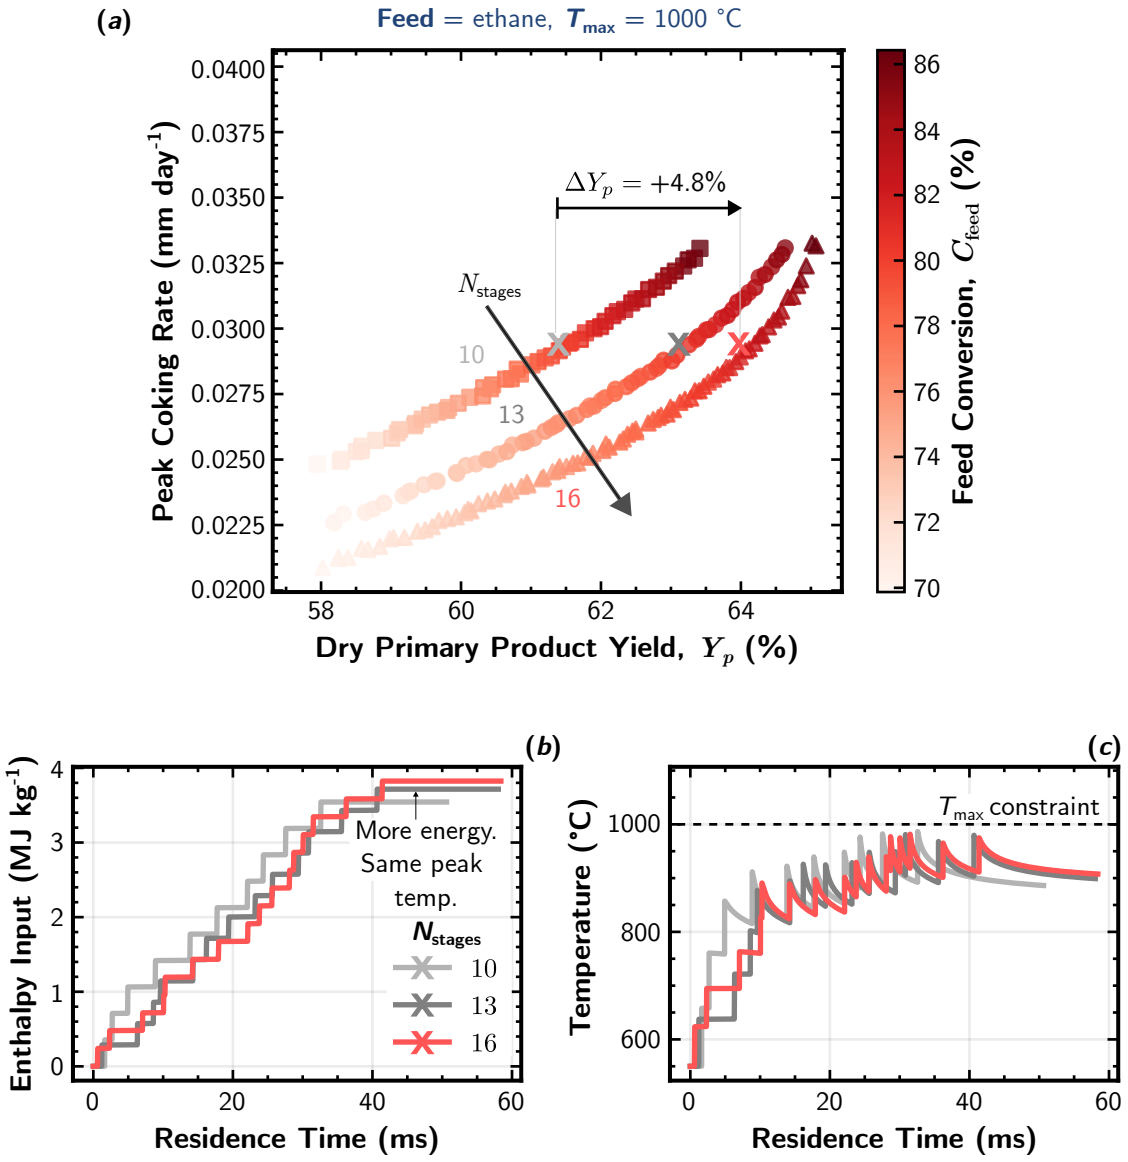


Figure 4.2: For an increasing number of stages $N_{\text{stages}} \in [10, 16]$: **(*a*)** the Pareto front of reaction performance **(*b*)** the cumulative work input distribution and **(*c*)** the temperature distribution

The results presented here imply that the shockwave reactor (SWR) introduced in Sec. 2.2.1 would have a reaction performance worse than the turbo-reactor. This is because in the SWR, all the energy is supplied to the reaction in one go rather than being distributed over a series of stages. This controllability and modularity of the turbo-reactor is a key advantage of the concept.

For a fixed coking rate with an ethane feedstock, increasing the number of stages from 10

to 16 enables a relative increase in yield and feed conversion of 4.8% and 8.2% respectively (see Fig. 4.2(a)). These improvements are significant, but diminish for more than 16 stages. Figures 4.2(a) & 4.2(b) suggest that for a fixed coking rate, more energy can be added for the same peak temperature, leading to the observed yield and conversion improvements.

Pareto Fronts: Peak-Temperature Constraint

For n-hexane feedstock, the maximum allowable temperature constraint is repeatedly relaxed from 940 to 1000 °C. Lower temperatures were tested, but these cases did not meet the predefined constraints. Figure 4.3 shows that for a fixed coking rate, only a 0.47% relative increase in the primary product yield is possible within the temperature range. This suggests that the theoretical maximum yield has been reached and that higher temperatures do not provide sizeable advantages. Interestingly, for a fixed yield, a higher peak temperature constraint enables lower coking, but only marginally.

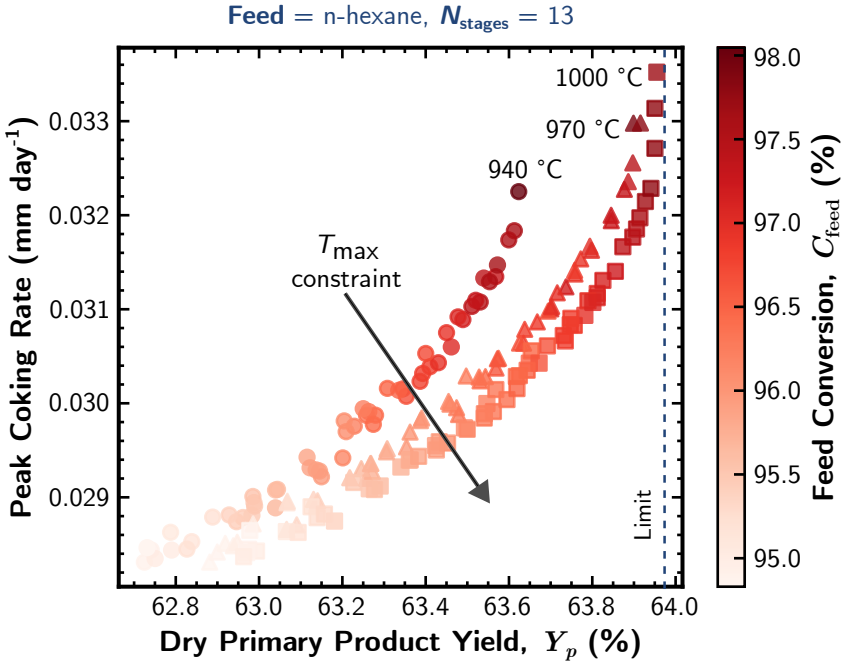


Figure 4.3: Effect of increasing the maximum allowable temperature level constraint on the Pareto front of reaction performance

Pareto Fronts: Non-Uniform Enthalpy Input Per Stage

Up to this point, the level of energy supplied to the reaction was optimised uniformly for all stages, and the rate of energy supply to the reaction was fixed. This subsection expands the parameter space (*i.e.,* to 39 tunable parameters), allowing both the level and the rate of the energy supply to be optimised for each stage (see Fig. 4.1). For a 13-stage machine with an ethane feed, Fig. 4.4(a) demonstrates that a non-uniform energy supply

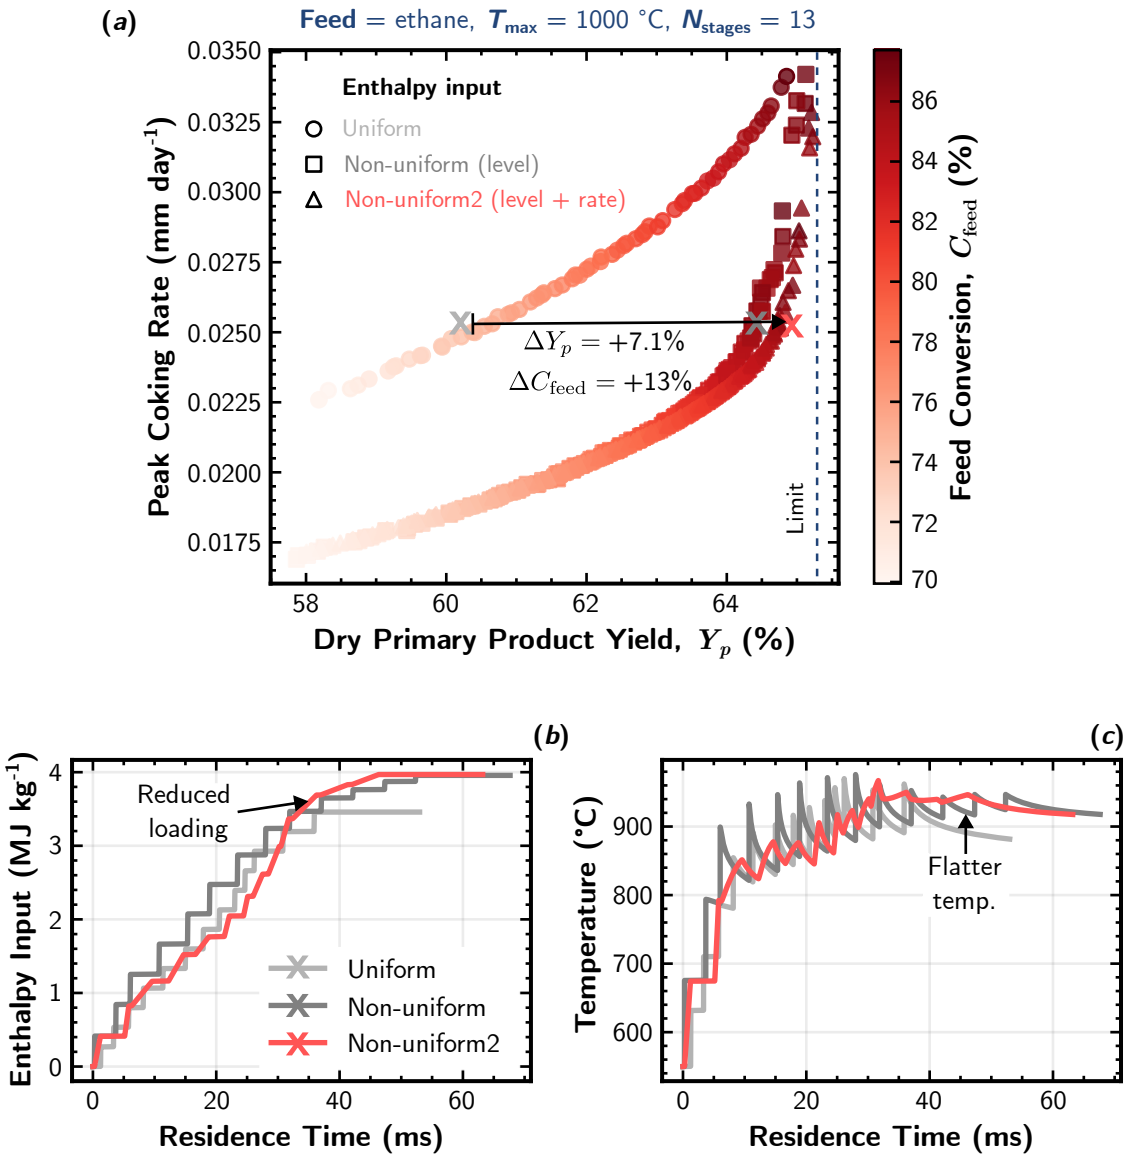


Figure 4.4: For uniform and non-uniform enthalpy input cases: **(*a*)** reaction performance Pareto fronts **(*b*)** cumulative work input distribution and **(*c*)** temperature distribution

has a notable impact on the reaction performance. For a fixed peak coke formation rate shown in Fig. 4.4(a), the relative increase in product yield and feed conversion is 7.1%

and 13% respectively. Adding an extra layer of complexity by enabling a variable ramp rate for the energy supply per stage has negligible benefits, indicating that the reaction yield has plateaued. Furthermore, the impact of the non-uniform energy supply level is appreciable only for low-coking-rate designs. If coke deposition can be mitigated using cooled blading, then designs operating towards the top right of Fig. 4.4(a) are possible, and there is no benefit from non-uniform energy supply.

Figures 4.4(b) & (c) suggest that the observed improvement in overall reaction performance is due to lower energy input rates towards reaction completion, flattening the temperature profile. Lower loading in the hot-end of the machine may also be beneficial for the aeromechanical design; however, the additional capital cost of a variable rotor design may not be justified.

4.1.2 Summary

It has been demonstrated in this section that disturbances from interstage intense turbulent dissipation processes are filtered out and suppressed by the highly accelerating stator blade row. Therefore, it has been shown that the level of specific energy input is independent of the interstage vaneless space length down to less than one diffuser chord length. As a result, the designer is free to redistribute the energy supply to the reaction to improve the reaction efficiency. To do this a constrained multi-objective optimisation investigation was performed. From this study, the following conclusions have been drawn.

1. A non-uniform multi-stage arrangement could be used to design a reaction-efficient temperature profile with a 10% yield improvement relative to the datum uniform multi-stage arrangement.

2. By increasing the number of stages from 10 to 16, a 4.8% relative improvement in yield is possible. However, these improvements must offset additional capital costs to be financially viable.

3. For an n-hexane feedstock, relaxing the temperature constraint from 940 °C to 1000 °C provided only marginal improvements to the Pareto front.

4. The most substantial improvements in reaction efficiency have been derived from allowing an independent energy supply level for each stage, gaining 7% and 13% relative improvement in yield and feed conversion, respectively.

There is one caveat to consider. There is uncertainty in the accuracy of the kinetic model and the coking model. To incorporate a more accurate and hence costly kinetic model, a chemistry acceleration technique is needed (which will be developed in Chapters 6 & 7). However, qualitative trends should be indicative and provide useful information to make basic design decisions.

Most importantly, from the investigations presented here, it has been demonstrated that there is considerable value in tailoring the aerodynamic design to improve the reaction performance. However, when operated in the real machine, there will be a large divergence from the optimum reaction profile due to three-dimensional effects and aerochemical interactions. Therefore, to achieve the potential yield and feed conversion levels demonstrated here, a fast 3-D reacting CFD flow solver is still necessary.

4.2 Robustness of the Concept to Feed Variability

A second important feature of robustness of the aerodynamic concept is the feed and reaction flexibility. In the context of endothermic chemical reaction applications, a low-order numerical approach is used to systematically analyse the effectiveness of a universal stage design strategy for a broad range of different applications (*e.g.,* steam methane reforming, ammonia decomposition, etc.), chemical feedstocks (*e.g.,* methane, ethane, n-hexane, ammonia, etc.) and reaction states. If a universal stage design philosophy can be used for different chemical processes, this will significantly simplify the design cycle and reduce capital costs. However, since fluid properties—and hence Mach number, Reynolds number, etc.—vary substantially, this is a challenging task. Fortunately, the concept is robust and therefore can supply high levels of specific enthalpy input over a wide range of regimes.

Pages 93–109:

Decarbonisation of High-Temperature Endothermic Chemical Reaction Processes Using a Novel Turbomachine: Robustness of the Concept to Feed Variability

(*J. Glob. Power Propuls. Soc.* 2024)

GPPS | JOURNAL

JOURNAL OF THE GLOBAL POWER AND PROPULSION SOCIETY
journal.gpps.global

Decarbonisation of high-temperature endothermic chemical reaction processes using a novel turbomachine: robustness of the concept to feed variability

Original article

Article history:
Submission date: 14 November 2023
Acceptance date: 4 March 2024
Publication date: 30 April 2024

This is the updated version of a paper originally presented at the Global Power and Propulsion Technical Conference, GPPS Hong Kong23, October 17–19, 2023.

Check for updates

***Correspondence:**
DR: dylan.rubini@eng.ox.ac.uk

Peer review:
Single blind

Copyright:
© 2024 Rubini et al. This is an open access article distributed under the Creative Commons Attribution Non Commercial No Derivatives License (CC BY-NC-ND 4.0). Unrestricted use, distribution, and reproduction of the original work are permitted for noncommercial purposes only, provided it is properly cited and its authors credited. No derivative of this work may be distributed.

Keywords:
compressor aerodynamics; chemical reactions; computational fluid dynamics; turbo-reactor; rotordynamic reactor; turbomachinery design

Citation:
Rubini D., Karefyllidis N., Rosic B., Xu L., and Nauha E. (2024). Decarbonisation of high-temperature endothermic chemical reaction processes using a novel turbomachine: robustness of the concept to feed variability. Journal of the Global Power and Propulsion Society. 8: 111–126. https://doi.org/10.33737/jgpps/185623

OPEN ACCESS

Dylan Rubini[1,*], Nikolas Karefyllidis[1], Budimir Rosic[1], Liping Xu[2], Elina Nauha[3]

[1]*Oxford Thermofluids Institute, University of Oxford, Oxford OX1 3PJ, UK*
[2]*Whittle Laboratory, University of Cambridge, Cambridge CB3 0DY, UK*
[3]*Coolbrook Oy, Pieni Roobertinkatu 9 00130, Helsinki, Finland*

Abstract

This paper presents a revolutionary turbomachinery concept, referred to as the turbo-reactor, which has the potential to replace gas-fired radiant furnaces and decarbonise a wide range of hard-to-abate, high-temperature endothermic chemical reaction processes. Although previous studies by the authors have confirmed the feasibility of using a turbo-reactor for steam cracking reactions, the numerical investigation presented in this work broadens the scope of potential applications for the machine to a variety of energy-intensive chemical processes, including those used for hydrogen production. This step change in technology could be the catalyst needed to enable rapid scale-up of low-carbon hydrogen technology. The innovative design of the turbo-reactor is fundamentally based on converting all of the imparted mechanical energy into internal energy, rather than pressure. This enables temperatures of up to 1,700°C to be achieved within an axial length on the order of one metre, resulting in an increase in the power density of 50 to 1,000 times compared to a surface heat exchanger. This paper presents the first comprehensive analysis of the turbo-reactor's robustness and controllability across a broad spectrum of feeds, chemical reaction stages, Mach number regimes, and operating points, conclusively demonstrating the feasibility of a universal stage design strategy for repeatedly imparting and dissipating energy for various endothermic reaction processes.

Introduction

Motivation and background

Decarbonisation of hard-to-abate high-temperature endothermic chemical reaction processes, such as steam cracking of hydrocarbons, ethylene dichloride cracking, plastic recycling, hydrogen production through steam methane reformation (SMR), and ammonia decomposition, is a significant challenge of this decade (Thiel and Stark, 2021). These processes are essential components of large-scale global industries. For example, the hydrocarbon cracking market was recently valued at $240 billion, with a global annual capacity of 400 million tons per annum (Mtpa) (Bender, 2014). This is expected to grow to $350 billion, with a capacity of 600 Mtpa, by 2030. Furthermore, with the rapidly growing support for low-carbon hydrogen and ammonia for transport, industry, energy storage, and power generation, it is expected that their demand will increase exponentially, potentially reaching to 200 Mtpa by 2030, with a market value of $320 billion (IEA, 2022).

The primary challenge is that large-scale energy-intensive sectors (e.g., chemical/petrochemicals steel, and cement) account for almost 20% of global CO_2 emissions (IEA, 2021), mainly due to heat generated by fossil fuel combustion. To address this challenge, a new class of turbomachines has been developed to replace gas-fired surface heat exchangers (specifically the firebox/furnace unit) used in these industrial processes, as shown in Figure 1.

Heavy industry can be transformed radically by exploiting the complex flow physics and controllable energy conversion processes possible within a turbomachine. This novel turbomachine, known as the RotoDynamic Reactor (RDR) or turbo-reactor, works on the principle of directly imparting mechanical energy to the working fluid supplied by a renewably powered electric motor as the mechanical driver (Rubini et al., 2021; Coolbrook Oy, 2022). As a result, Figure 1 demonstrates that the turbo-reactor can significantly increase power density by up to two orders of magnitude compared to a surface heat exchanger, leading to a reduced plant footprint, as well as lower capital, operating, and maintenance costs. This is achieved by converting almost all of the imparted mechanical energy into internal energy, rather than compressing the gas as is done in a compressor (which is another example of an energy-imparting machine). Therefore, high working fluid temperatures of up to 1700°C can be achieved within short axial distances of less than one metre, depending on the working fluid properties. This can result in a significant improvement in the efficiency and yield of the reaction (Rubini et al., 2022b). Previous work by the authors presents a detailed introduction and analysis of stage design, energy transformation train, and fundamental flow physics (Rubini et al., 2021, 2022a,b; Karefyllidis et al., 2023).

Figures 2 and 3 highlight the significant role of the RDR in decarbonising a diverse set of high-temperature endothermic reaction processes for the production of high-value commodity chemicals in the chemical industry, such as steam cracking, thermochemical plastic waste recycling, SMR for hydrogen production, and ammonia decomposition for hydrogen recovery. Alternative strategies to decarbonise these sectors include hydrogen-fired (Weydahl et al., 2013) or electric furnaces (Delikonstantis et al., 2019), but these suffer from large thermal resistances that hinder energy transfer (Venkataraman et al., 2003), low power density, and non-uniform temperature profiles within the tubes, all of which can degrade the efficiency of the process. Furthermore, there are additional

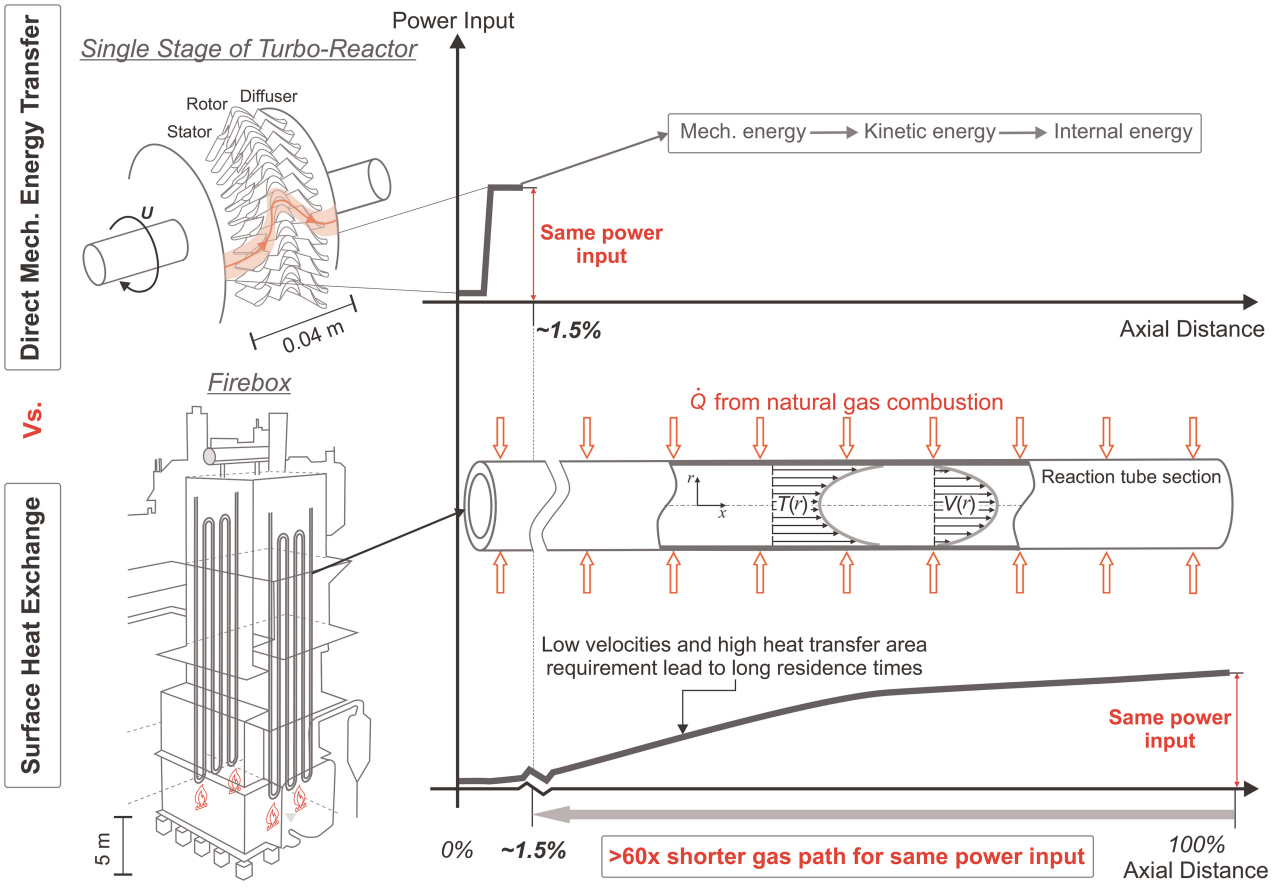


Figure 1. A schematic analysis of the energy transfer mechanism in a surface heat exchanger against that in the turbo-reactor.

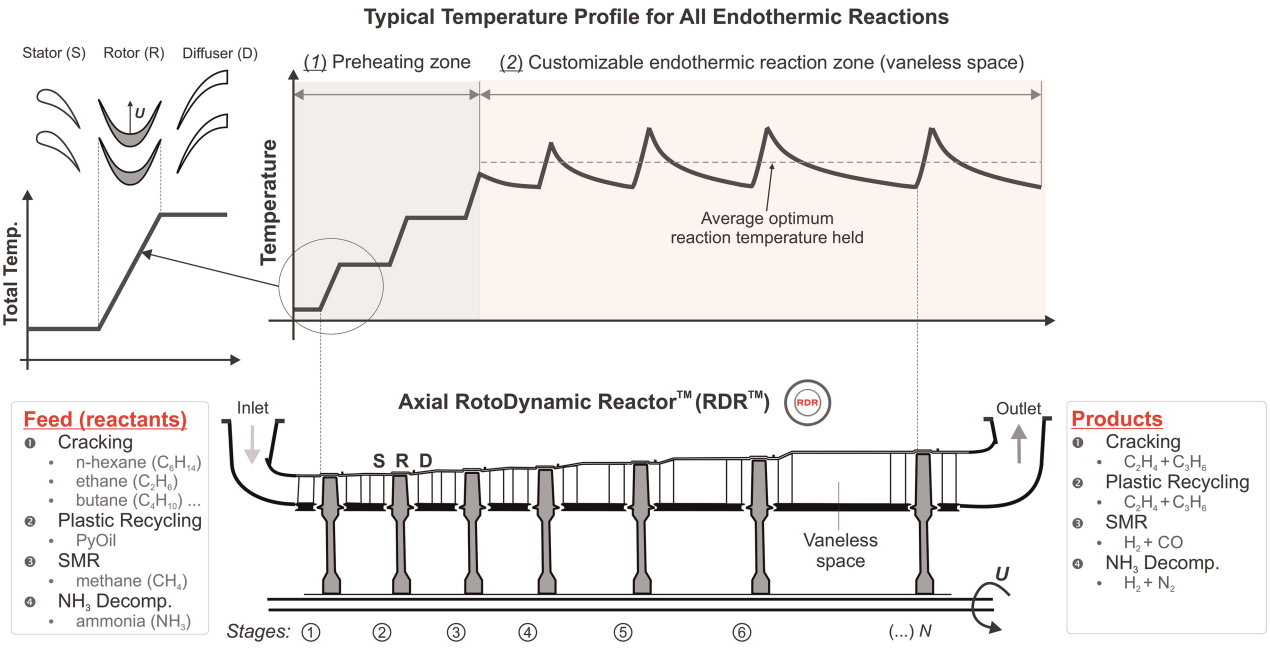


Figure 2. A schematic of the elemental stage and multistage environment for the RotoDynamic Reactor, illustrating a typical temperature profile for some example endothermic chemical processes that could be driven by the RDR.

challenges related to scalability at high temperatures for electric furnaces and low energy conversion efficiency $\leq 60\%$ (Lu et al., 2022) and nitrous oxide emissions for hydrogen-fired furnaces.

Objectives of this study

This first-of-its-kind paper demonstrates the applicability of the RDR for decarbonising a wide range of endothermic reactions within the chemical industry (see Figure 3). This study confirms the robustness and controllability of the system subjected to different feedstocks, variable reaction states, and a range of operating points. A broad spectrum of operating states is generated by harnessing variations in thermophysical/thermochemical fluid properties corresponding to different feedstocks (e.g., n-hexane, methane, etc.) and reaction progress states. The RDR is shown to be resilient and flexible, and energy can be effectively imparted and dissipated despite using a fixed stage design and boundary conditions.

The turbo-reactor: direct energy transfer for endothermic chemical reactions

Controllable elemental stage design for ultra-high power density energy input

The turbo-reactor concept and detailed design have been developed over many years of collaboration with Coolbrook Oy (2022) and are based on the original patent of Bushuev (2016). This new machine unlocks a radically new design space for high-power-density energy-imparting machines, as shown in Figures 2 and 4. Figure 4 compares the single-stage constant speed $\phi - \psi$ and $\phi - C_{\psi_{TS}}$ characteristics between the turbo-reactor and an axial compressor. As the primary objective of an axial compressor is the increase of pressure, the work coefficient ψ is limited by isentropic efficiency constraints throughout the operating range (see Figure 4b), and system instability considerations at low flow coefficients ϕ (see Figure 4c). For the RDR, by eliminating the requirement for an increase in static pressure, these constraints are bypassed within the relevant operating range (that is, $p_{exit}/p_{0,in} \leq 1$). This results in the following key design differences from a compressor:

1. For low molecular weight fluids, an ultra-high loading coefficient of up to 8.0 is possible (see Figure 4b). This is more an order of magnitude higher than that of a typical axial compressor, minimising the number of stages required for heating. Figure 4b indicates that the work coefficient increases approximately linearly with the flow coefficient, before turning to the vertical due to feed-composition-dependent choking limitations.
2. Within the choking limits of the working fluid, the RDR can leverage very high relative Mach numbers $\phi \approx 1.0$ and flow coefficients (up to $\phi \approx 1.0$) without the risk of system instability and overall performance

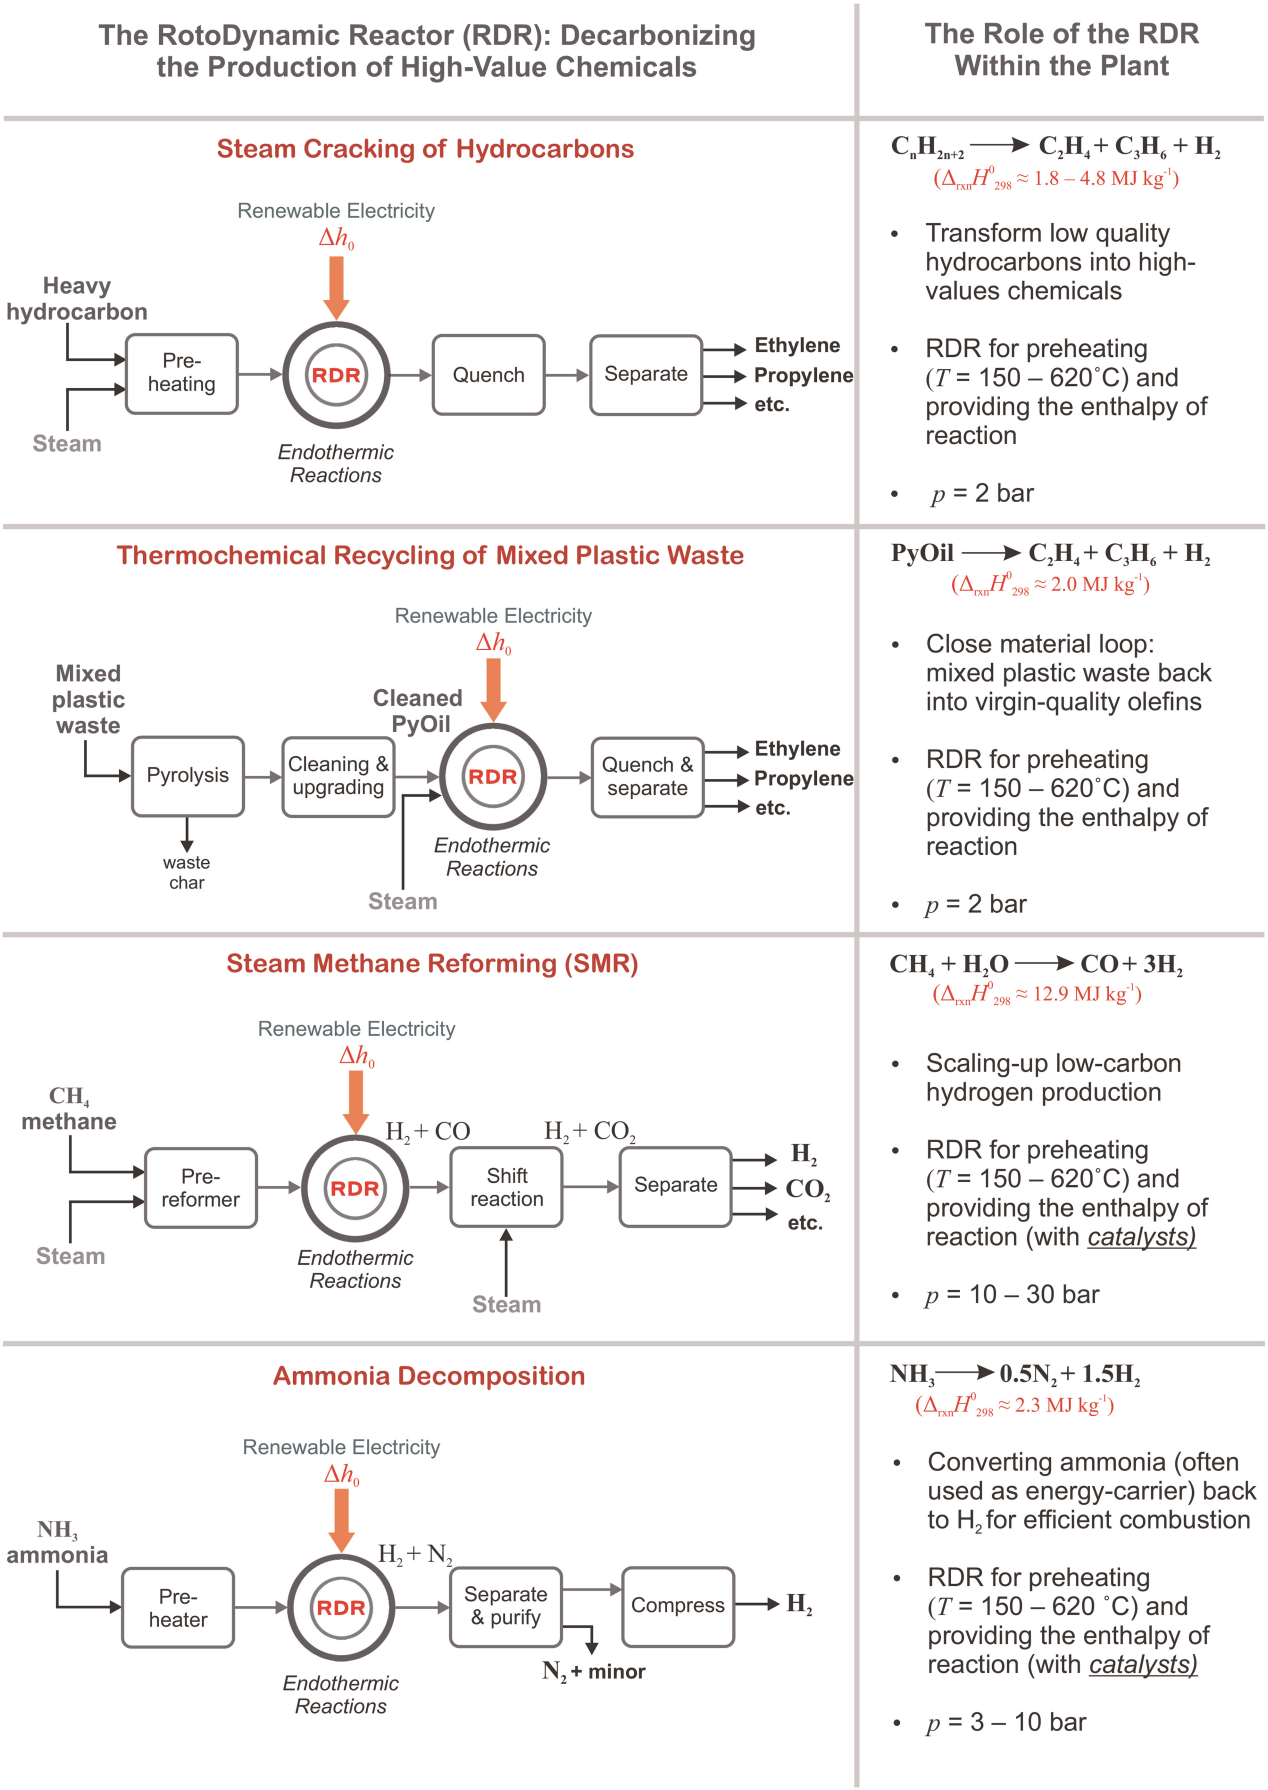


Figure 3. Schematic illustration of potential opportunities for turbomachines to decarbonise high-temperature endothermic reaction processes in the chemical and petrochemical industry.

degradation. This enables a high flow capacity per unit frontal area, resulting in a smaller machine footprint.

3. Since there is no requirement to support a global adverse pressure gradient, Figure 4c illustrates that the onset condition for violent dynamic system instability is circumvented as $(\partial C_{\psi_{TS}}/\partial\phi) < 0$ (Greitzer et al., 2007). Therefore, the RDR is robust, controllable, and inherently free of large-scale system instability (i.e., negative aerodynamic damping) over a wide range of operating states. However, it is noted that local instability may still grow within certain blade rows.

These features are achieved by imparting almost all the available exergy into the fluid as kinetic energy (KE) using high-turning supersonic impulse rotor blades (see Figure 2) rather than pressurising the gas. Through aerodynamic losses, this KE is then rapidly dissipated into internal energy within the diffuser and vaneless space. This allows very high temperatures to be attained in a fraction of a millisecond without a pressure rise across the stage. For example, a temperature change greater than 300 °C per stage can be achieved with air as the working fluid. As a result, the flow path length required to heat a given mass of fluid to a specified temperature can be reduced by nearly two orders of magnitude compared to a furnace (see Figure 1).

In a multistage environment, the flow must be reconditioned before the next rotor row. A stator vane is used to rapidly accelerate the flow to help it recover before entering the downstream row and to provide the necessary high Mach number with a large negative swirl angle (to the rotor). The favourable pressure gradient attenuates residual flow non-uniformities from the upstream mixing process, ensuring high performance across the multistage architecture (Karefyllidis et al., 2023).

Addressing the limitations of surface heat exchange

Compared to radiant furnaces, the RDR provides a fundamental advantage: mechanical energy is directly transferred into the fluid, making it the hottest part of the system. This is in contrast to furnaces, where the walls are the hottest part of the system due to the heat transfer process. This inherent feature of volumetric energy addition bypasses two key challenges encountered in radiant coils. First, the heat transfer area requirement is avoided, allowing a significantly shorter gas path (see Figure 1). Reaction selectivity is a function of residence time; therefore, the primary product yield can be increased.

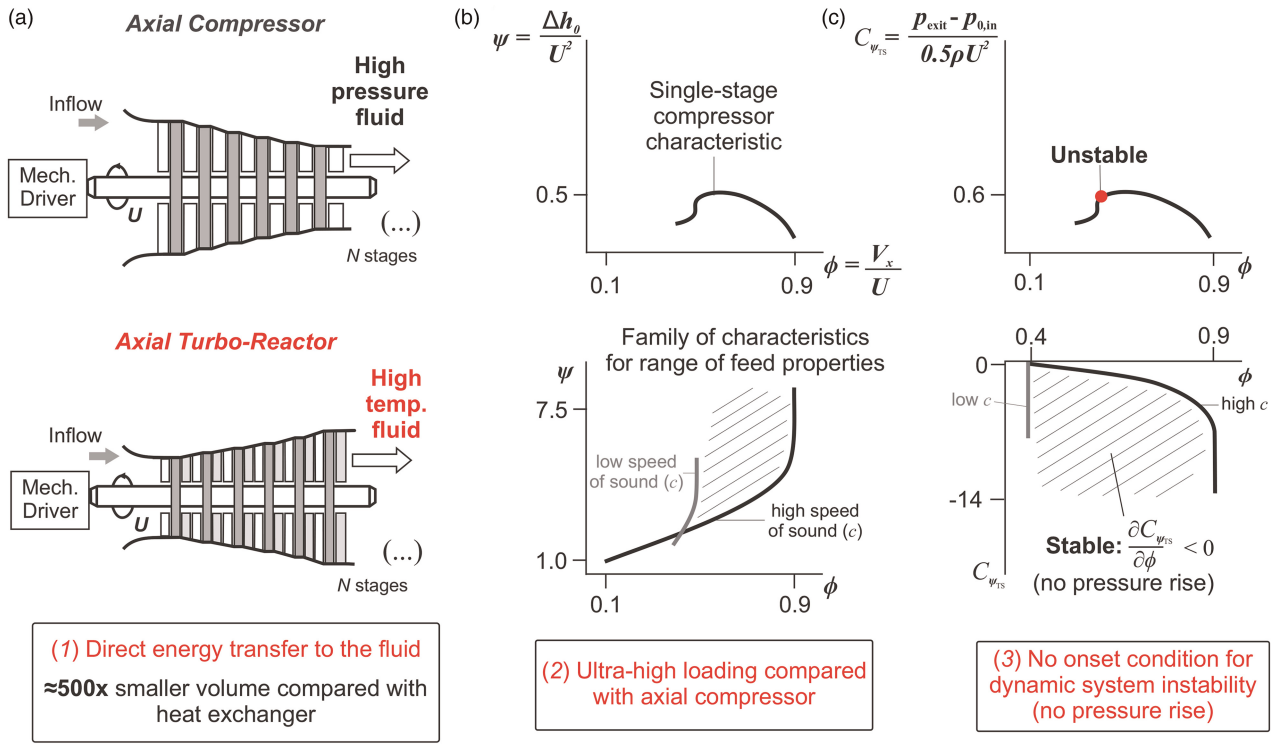


Figure 4. A comparison between an axial compressor and turbo-reactor (both energy imparting machines), showing (a) the primary objectives, (b) the single-stage flow coefficient ϕ against stage loading ψ characteristic speedline, and (c) the single-stage flow coefficient ϕ against total-to-static pressure rise coefficient $C_{\psi_{TS}}$ characteristic.

Second, as indicated in Figure 1, the thermal and velocity boundary layers within the tubes are thick and mismatched. As a result, the majority of the fluid's mass is concentrated at the centre, but the highest temperature is at the surface. This results in coke deposition on the walls due to high metal temperatures (for hydrocarbon feeds) and overcracking due to nonhomogeneous reaction progress. By avoiding these limitations, the RDR enables a more homogeneous and efficient reaction for all chemical processes, as well as lower coking rates for hydrocarbon-based feeds.

Fine-tuned control of the endothermic reaction dynamics

For high-temperature endothermic chemical reactions, where the enthalpy of reaction $\Delta_{\mathrm{rxn}}H^0_{298}$ is positive, a turbomachine can be used to replace conventional radiant furnaces. The turbo-reactor preheats the gas mixture to a temperature where the reaction proceeds efficiently, followed by further customised stages to maintain the average optimum reaction temperature during the subsequent enthalpy-draining reaction, as shown in Figure 2. This paper will show the versatility of the RDR using four example applications (see Figure 3) summarised below.

1. Steam cracking: transforming heavy hydrocarbons into higher-value olefins such as ethylene (C_2H_4) and propylene (C_2H_6). This is a thermal decomposition (i.e., pyrolysis) reaction in an oxygen-free environment.
2. Thermochemical recycling of mixed plastic waste (Kusenberg et al., 2022): decomposing pyrolysis oil (PyOil) back into virgin-quality olefins (see Figure 3). PyOil is converted to olefins using the same process mentioned above (see 1.).
3. Steam methane reforming: the RDR is used to produce syngas ($CO + H_2$) from natural gas (CH_4). This syngas is processed and separated into hydrogen (H_2) and carbon dioxide (CO_2) (Rostrup-Nielsen and Christiansen, 2011).
4. Ammonia (NH_3) decomposition: converting ammonia into hydrogen (H_2), which is becoming an important process to release H_2 from an ammonia energy carrier for cleaner and more efficient combustion (Jackson et al., 2019).

A common feature among the working fluids used in these applications is that they are all superheated vapours following the ideal gas law with a temperature-dependent isobaric heat capacity (i.e., the mixture is a semi-perfect gas). In the first three applications, superheated steam is added to the feed to reduce the coke (i.e., solid carbon) deposition on the surfaces. Therefore, steam plays a crucial role in determining the properties of the mixture. The fact that the mixture is a single-phase, semi-perfect gas (rather than a non-ideal fluid) significantly simplifies turbomachinery design and modelling processes.

The objective of the non-reacting, ultra-fast preheating section ($150\ ^\circ\mathrm{C} \leq T \leq 620\ ^\circ\mathrm{C}$) for all chemical processes is to maximise the flow of energy into the fluid within a minimum distance and time scale. In some configurations of the preheating zone, the vaneless space can be eliminated by integrating the diffuser and the stator. As a result, the flow can enter successive stages with a higher and higher flow coefficient, enabling a potential increase in stage loading across the multistage architecture (until an equilibrium is reached).

Once preheating is complete and the gas mixture is at the reaction activation temperature, the vapour enters the reaction zone where the turbo-reactor provides a tailored temperature profile based on the application-specific reaction dynamics (see Figure 2). This is achieved using a non-uniform axial distribution for the interstage vaneless space length (see Figure 2). For each stage, the axial dimension depends on the local heat of reaction, which is a function of the thermochemical state.

Numerical methodology

In the following sections, the versatility of the turbo-reactor will be demonstrated in its ability to operate effectively across a wide range of feedstocks and reaction conditions for various chemical processes powered by the RDR. The basic requirements for each process, based on the design of the datum stage, are summarised in Table 1. To evaluate the performance of the machine under these varying conditions without relying on computationally intensive reacting flow simulations, a framework has been developed to decouple the aerodynamics and chemical kinetics. This is introduced in the next section.

A decoupled framework for evaluating the effects of chemistry on aerothermal performance

A low-order approach is taken to decouple computational fluid dynamic (CFD) simulations and chemical kinetics. This is summarised in Figure 5 and can be explained as follows. Offline, 1-D detailed kinetic simulations are conducted for steam cracking (of an n-hexane and ethane feed), SMR, and ammonia decomposition

Table 1. Approximate preliminary requirements for the preheating and reaction zones based on the datum stage.

Feed/ process	Preheating zone (150°C < T < 900°C)			Reaction zone ($T \approx$ 900°C)		
	Approx. no. stages	Residence time (ms)	Enrgy-input (MJ kg^{-1})	Approx. no. stages	Residence time (ms)	Energy-input (MJ kg^{-1})
n-Hexane	8	4.0	2.4	6	11	1.8
Ethane	8	4.0	2.4	13	27	3.9
SMR: methane[a]	7	3.5	2.1	12–18	20–250	4.0–12.0
Ammonia[b]	8	4.0	2.4	12	170	3.6

[a]With Nickel catalyst[c]
[b]With Ruthenium catalyst[c].
[c]Surface-area-to-volume ratio = 400 m^{2}m^{-3}.

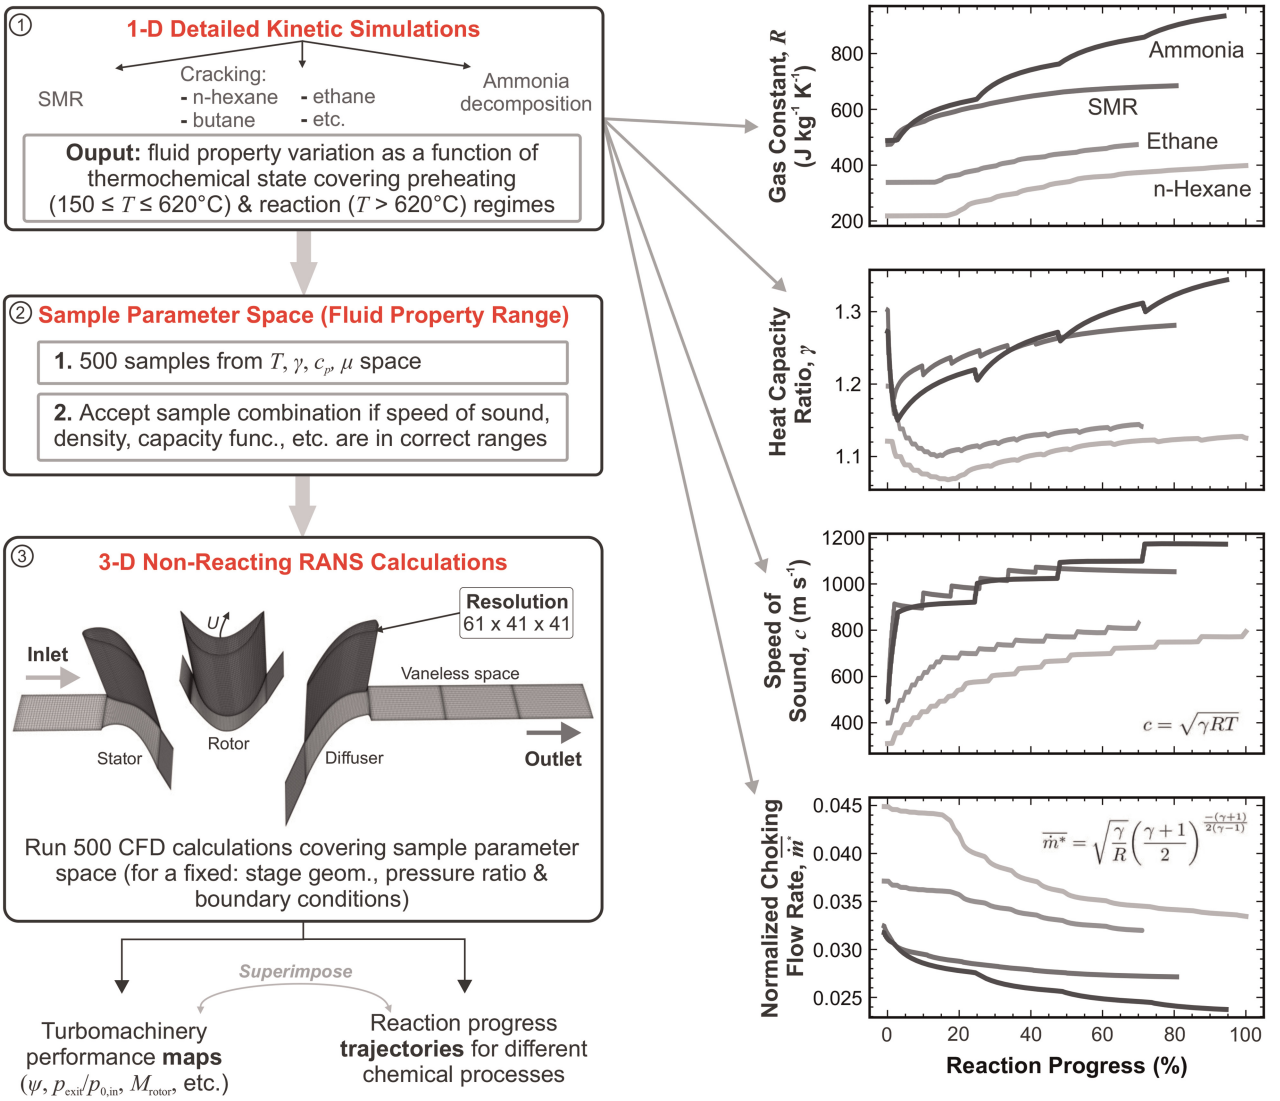


Figure 5. (Left) The methodology for constructing performance maps of the RDR for a broad range of chemical processes and reaction states and (Right) temperature-dependent variations in thermophysical properties, speed of sound, and non-dimensional choking mass flow rate w.r.t. reaction progress in the range 150–1,000°C.

reactions across a broad temperature range (150°C $\leq T \leq$ 1,000°C) covering both the preheating and reaction zones. All of the reactions are taken near completion to capture a broad range of reaction states.

The impact of the chemical reaction is effectively lumped into changes in the fluid properties, with four variables forming the parameter space: static temperature T, heat capacity ratio γ, isobaric specific heat capacity c_p and dynamic viscosity μ. By sampling from this fluid property space (see Figure 5) and performing a non-reacting CFD calculation for each sample, the performance can be evaluated under different reaction conditions without performing costly reacting flow simulations. The aerothermal performance of the turbomachine is characterised in the same way as conventional turbomachines using the non-dimensional groups shown in Equation 1 and defined in Dixon and Hall (2013):

Total pressure ratio

Reynolds number

$$\frac{p_{02}}{p_{01}}, \frac{\Delta T_0}{T_{01}}, \psi = f\left(\frac{\dot{m}\sqrt{\gamma R T_{01}}}{D^2 p_{01}}, \frac{\Omega D}{\sqrt{\gamma R T_{01}}}, Re, \gamma\right) \quad (1)$$

Stage loading coeff.

Capacity func.

Non-dim. blade speed

where p_0 is the total pressure, R is the gas constant, $\dot{m}$ is the mass flow rate, D is the mean diameter and Ω is the rotational speed. The terms highlighted in red are composition-dependent, indicating that turbomachinery performance is a direct function of the state of progress through each reaction. It is noted that there are two key implicit assumptions in this analysis. First, it is assumed that the aerodynamics and chemistry can be decoupled, and therefore are not assumed to be strongly interacting. Second, the fluid properties are assumed to be constant across a stage. In general, these are crude assumptions; however, these approximations allow us to gain rapid insight into how different reaction states change aerothermal performance.

Block (1): 1-D chemical kinetic modelling

For each feedstock and chemical process illustrated in Block (1) of Figure 5, a chemical reaction mechanism is automatically generated using RMG-PY (Gao et al., 2016). This includes both gas-phase reactions and heterogeneous surface-phase (catalytic) reactions. An in-house 1-D reactor model, built on top of a plug-flow reactor (PFR) model and a stiff ordinary differential equation solver within CANTERA (Goodwin et al., 2022), is used to approximate the enthalpy supplied by the rotor separated by user-defined time intervals representing the residence time per stage. Both ammonia decomposition and SMR require catalysts. For both, the surface-area-to-volume ratio is set to 400 m^2 m^{-3}, which could be tuned if required.

For each reaction, the validity of the semi-perfect gas assumption was analysed. The Redlich–Kwong real gas equation of state was used to simulate the gas mixtures at the relevant operating points. For almost all regimes, the compressibility factor Z and fugacity coefficient Φ deviated by less than 2% from their ideal gas values. For SMR, at low temperatures of $T \leq$ 200°C and at high pressures of 30 bar, the deviation from ideality increased to a maximum of ≈8% for Z and Φ. However, the error quickly dropped to a negligible level once the temperature rose by 100°C. This provides confidence in the semi-perfect gas assumption over an extensive range of conditions and feedstocks.

Block (2): sampling the fluid property space

The parameter space is formed by randomly sampling T, γ, c_p, and μ. As the non-dimensional groups on the right-hand side of Equation 1 involve *combinations* of these parameters, such as the stagnation speed of sound $\sqrt{\gamma R T_{01}}$, only samples where the capacity function, speed of sound, non-dimensional blade speed, etc., fall within the known ranges, are accepted. These *known ranges* are determined from precursor 1-D chemical kinetic simulations covering a range of feedstocks and feed conversion states (see previous section).

Block (3): 3-D steady non-reacting RANS CFD modelling

To evaluate aerothermal performance, 500× steady three-dimensional, non-reacting RANS calculations using a single stage were performed for each sample in the parameter space (see Block (3) in Figure 5). This was achieved

using the flow solver TBLOCK (Denton, 1983, 1986, 1992) with a mixing-length turbulence treatment. The mixing length-scale is tuned on a block-basis and is set by the turbulent length scale calculated from high-fidelity LES. The computational setup, which employed a simplified computational domain and numerical treatment, allowed hundreds of CFD simulations to be run within a few hours. To assess mesh independence, three grades of mesh density were simulated. The mass-averaged flow coefficient, stage loading coefficient, and rotor exit Mach number error were within 0.7%, 1.9%, and 2.0%, respectively. This gives us confidence in the trends presented in the following sections.

Although the mesh is relatively coarse and the turbulence modelling relatively simple, it is emphasised that the main objective is to determine qualitative trends rather than absolute values. Moreover, the main quantity of interest is the bulk work coefficient, which is found to be less sensitive to turbulence modelling and mesh density for the impulse-type blades used here.

Robustness of the turbo-reactor concept to feed variability

This paper presents the first instance in turbomachinery in which the flow is reacting across the majority of the gas path length. This results in significant variations in the thermophysical fluid properties from inlet to outlet. These changes cover a broad range of molecular weights; therefore, leading to reaction-dependent variations on turbomachinery performance, more specifically, due to changes in the gas dynamics ($c = \sqrt{\gamma RT}$, γ, $\overline{\dot{m}^*}$, etc.). In this section, we demonstrate the controllability of the machine over a broad range of Reynolds numbers, Mach numbers, and flow coefficients to illustrate the robustness of the concept to feed variability.

Influence of Reynolds number on aerodynamic performance

The Reynolds number $Re = \rho V b/\mu$ (based on the stator exit conditions and true blade chord) varies with feed composition and reaction progress due to changes in fluid density and dynamic viscosity. Figure 6a suggests that stage loading is weakly dependent on Reynolds number over a wide range ($1 \times 10^4 \leq Re \leq 2.5 \times 10^6$), dropping by only 9% over the investigated range. As expected, more energy can be imparted into the fluid at high Reynolds numbers. Figure 6b indicates that this is because, in the low Reynolds number regime, there is a higher level of flow separation as the boundary layer is more sensitive and cannot negotiate local adverse pressure gradients. Higher flow separation leads to higher blockage and lower throughflow capacity, as well as increased flow deviation. Since stage loading scales with flow coefficient and flow turning, the specific energy input falls at low Reynolds numbers. However, the change is relatively minor (see Figure 6a).

Aerothermal performance maps covering a wide operating range

This section demonstrates the applicability of using a turbomachine for four different chemical processes using a uniform stage design, fixed annulus flow area, and fixed boundary conditions. These conditions include a fixed imposed pressure ratio $p_{exit}/p_{0,in} \approx 0.98$ and blade tip speed. As highlighted at the bottom right of Figure 7, by varying the fluid properties (see Figure 5) while fixing $p_{exit}/p_{0,in}$, the system naturally experiences variations in the operating point, set by the flow coefficient $\phi = V_x/U$. This variation is exploited to probe the response of the dynamical system (i.e., the RDR) subject to input perturbations (i.e., fluid properties). As a result, the

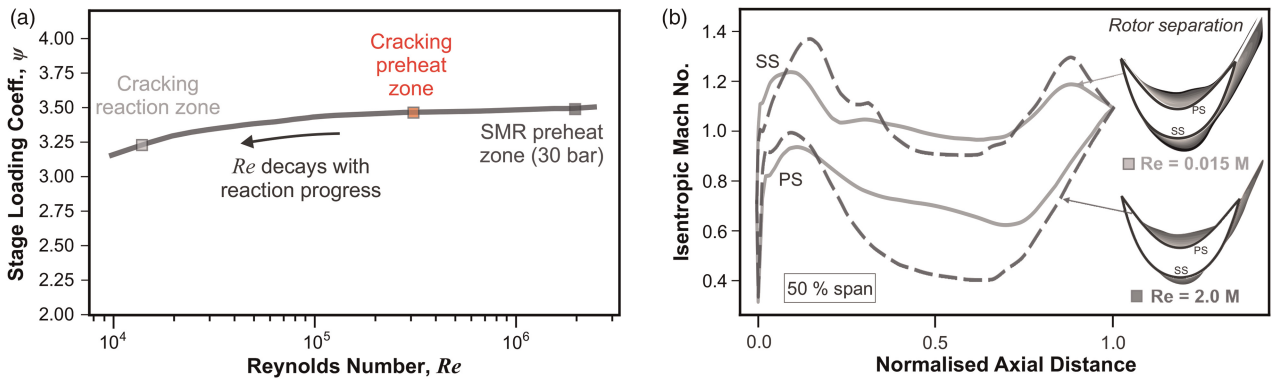


Figure 6. (a) Variation of stage loading coefficient with Reynolds number and (b) the isentropic Mach number distribution for two extreme Reynolds number states, corresponding to different reactions powered by the RDR.

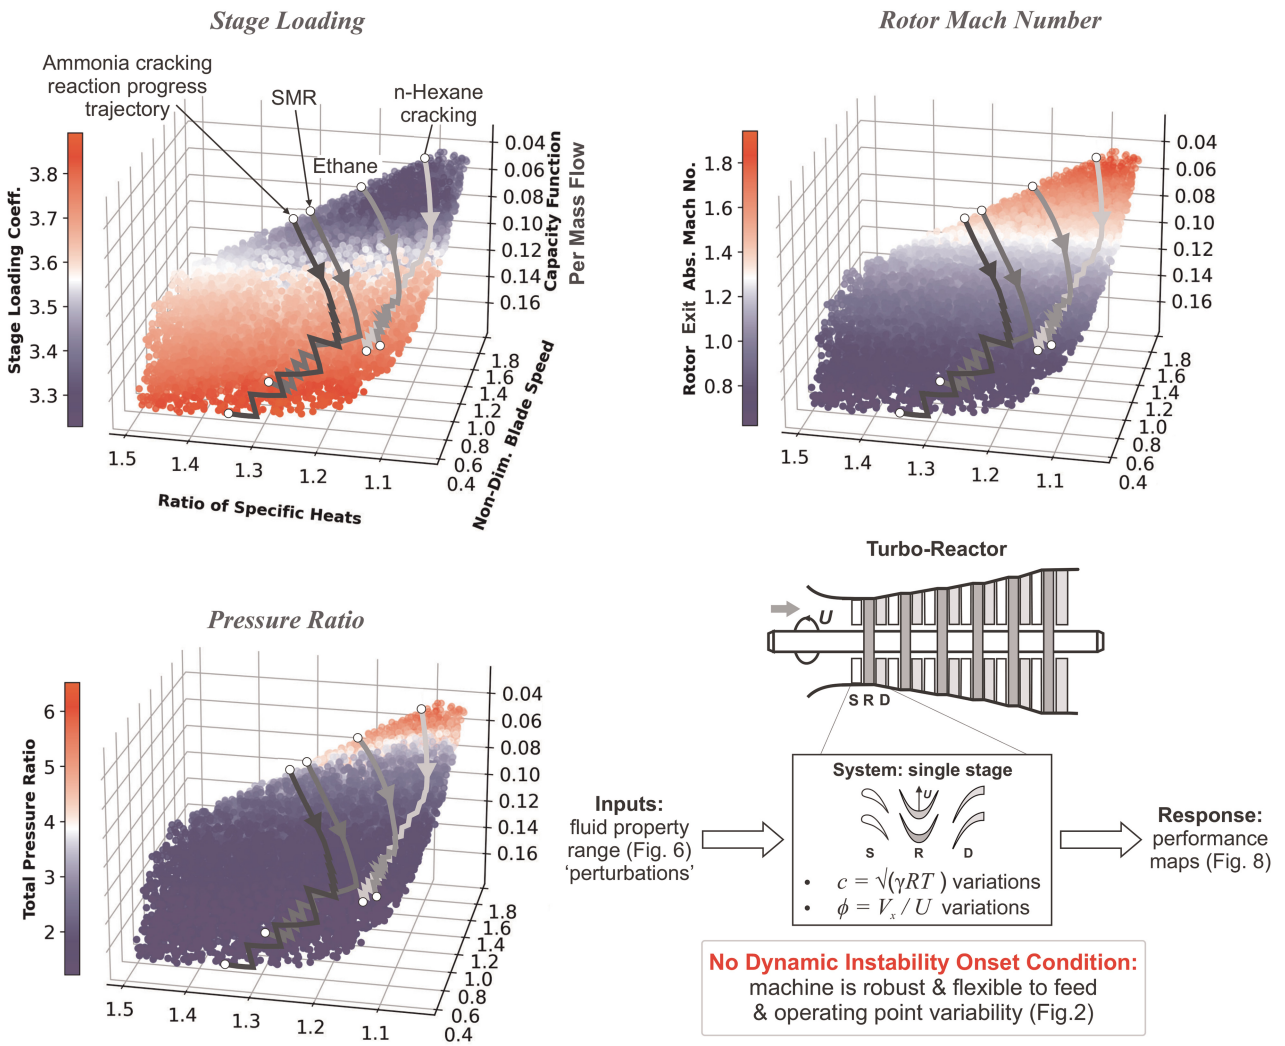


Figure 7. Aerothermal performance maps showcasing the reaction-progress trajectories with thick solid lines for four applications of the turbo-reactor (for a fixed *Re* = 100,000 and pressure ratio $p_{exit}/p_{0,in} = 0.98$).

robustness and controllability of the machine can be demonstrated over a wide range of fluid properties and hence operating regimes. Of course, during normal operation, the delivery pressure, stage design, and annulus flow area can be customised for each stage and for each chemical reaction to maintain and match the optimum velocity triangles.

Each thick solid line indicated in Figure 7 represents the trajectory through non-dimensional parameter space using decoupled detailed 1-D kinetic calculations. This detailed chemical analysis is performed independently, and the results are superimposed on the performance maps. These lines effectively illustrate how aerothermal performance varies along the streamwise direction across a multistage machine for each chemical process. For all reactions, the imposed pressure ratio is fixed; therefore, the Mach number decays across the multistage machine. Since density decreases and the flow area is fixed, this is accompanied by an increase in the flow coefficient toward the rear of the machine. Therefore, the stage loading is higher near reaction completion. At the beginning of the reaction, the feeds are initially separated in the non-dimensional parameter space (see Figure 7). However, as the reactions progress, they all converge toward the higher loading region.

More generally, two insights can be gained from the response of the performance metrics shown in Figure 7. First, the ratio of specific heats has a relatively small impact on the stage loading. Secondly, the dominant parameters driving changes in flow physics are the non-dimensional blade speed and flow coefficient. The non-dimensional blade speed is inversely proportional to the speed of sound (see Figure 5 and Equation 1) and varies by almost a factor of 4.5, whilst the flow coefficient, which sets the velocity triangles, varies by almost a factor of 2 (due to the density variation). The stage loading coefficient does not respond significantly, with a maximum deviation of 15%, illustrating the robustness of the turbo-reactor despite the use of a fixed elemental stage. This margin can be further reduced by adjusting the delivery pressure and modifying the flow area to match the flow

coefficient and the restore the velocity triangles. Since there is no onset condition for violent system instability (see Figure 4c), the machine can provide a high enthalpy input over a broad range of reaction and operating states.

Parameter space boundaries: extremes of the Mach number regime

Two extreme Mach number regimes, covering a wide range of feed molecular weights (2.0–86.0 g mol^{-1}) and hence the speeds of sound (280–1,150 ms^{-1}), are selected from the parameter space shown in Figure 7 and are compared in more detail in Figure 8. The low Mach number regime corresponds to an ammonia feed toward reaction completion, and the high Mach number regime corresponds to an n-hexane feed before the reaction starts. The first case represents a front stage of the machine, which is capacity-limiting due to the low speed of sound (see Figure 5 (*Right*)), and the second case represents a rear stage. The high Mach number case (see Figure 8 (*Right*)) is presented to illustrate the robustness of the flow at off-design regimes. However, in practice, the mismatched velocity triangles can be partially corrected by adjusting the delivery pressure to accommodate a higher mass flow rate, increasing the annulus flow area, and/or increasing the throat area to raise the swallowing capacity near the choke point.

Figure 8 provides three key insights into the flow physics of the system. First, within the rotor, the low Mach number case is mainly affected by PS separation and blockage, whereas for the high Mach number case, an increased positive incidence angle results in a higher level of SS separation contaminating the flowfield. This is because (I) the flow coefficient drops by almost a factor of 2 for the high Mach number case (II) the Mach number is supersonic, making the flow sensitive to small changes in incidence ($\leq$1.5°), and (III) a shockwave at the rotor inlet contributes to further SS boundary layer separation through shockwave boundary layer interaction (SWBLI).

Second, the gas mixture in the low Mach number regime (Figure 8 (*Left*)) consists of 70% hydrogen on a mole basis. As hydrogen has a low molecular weight and hence a high speed of sound, it can accommodate a higher flow coefficient at a given pressure ratio and a fixed annulus area. This allows for a greater change in tangential velocity whilst avoiding the SWBLI in the rotor, resulting in higher stage loading (see Figure 8). As the flow is far below the choking limit, an even higher flow coefficient can be achieved by increasing the delivery pressure (see Figure 4b). Therefore, it is evident that the local fluid properties (set by the local composition state)

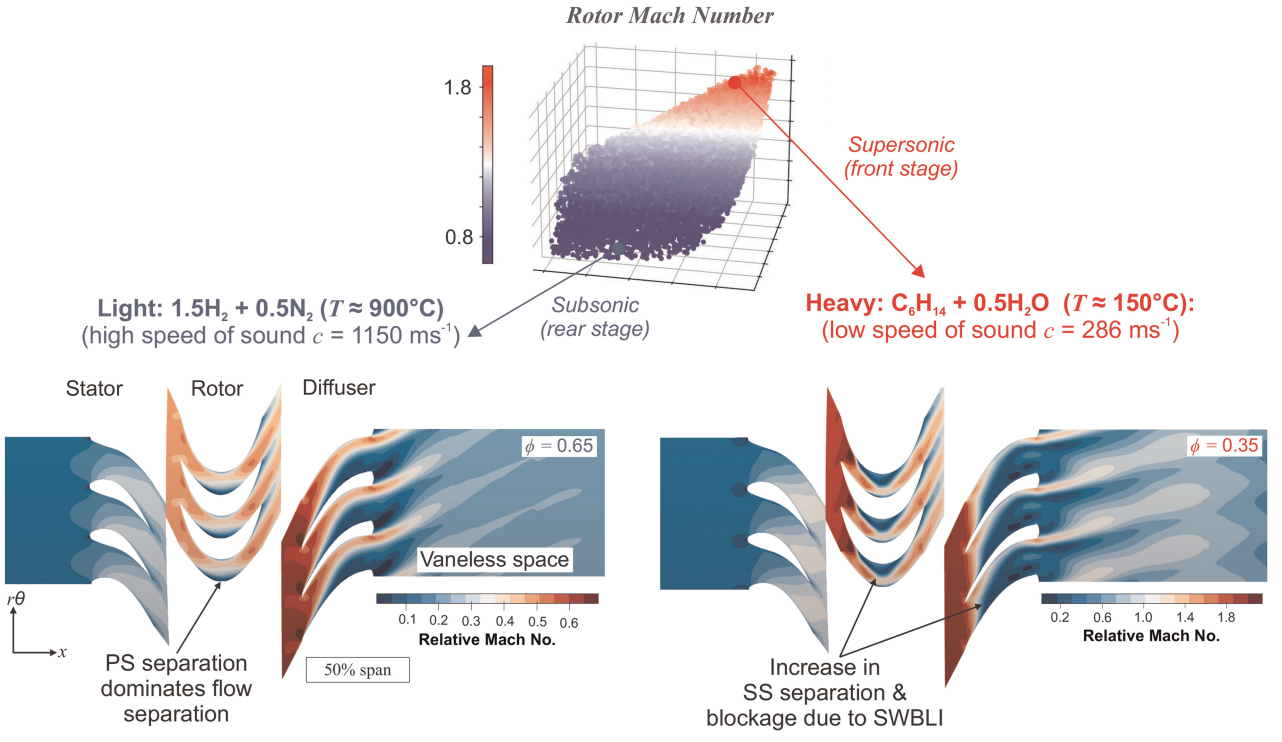


Figure 8. Comparison of the midspan relative Mach number distribution at two "extreme" corners of the parameter space where the fluid properties and local gas composition lead to two polar speed of sound regimes.

have a significant influence on the operating range, compressibility, and loading capability for a fixed geometry. With the current datum design, the operating range of the high molecular weight n-hexane feed (Figure 8 (*Right*)) is more restricted, and the $\phi - \psi$ characteristic is almost vertical (see the low speed of sound boundary in Figures 4b and 4c). However, for a given fluid, the design can be altered to balance the operating range and the specific energy input rate per stage.

Finally, the results presented demonstrate that the RDR can still achieve a reasonably high level of specific energy input ($\Delta h_0 \approx 220$ kJ kg^{-1}) despite large-scale flow separation due to incidence and SWBLI, as indicated in Figure 8 (*Right*). This flexibility of the machine is a direct consequence of the fact that isentropic efficiency does not impact the performance of the machine. Without a requirement to resist an adverse pressure gradient, entropy can be generated without dramatic repercussions for the specific energy that can be imparted into the fluid. Despite an elevated level of flow separation within the passage, the flow within the rotor is turned through an angle close to the blade metal angles. Therefore, the specific energy imparted into the fluid remains high, albeit at the expense of a reduced flow coefficient due to blockage. In the high Mach number regime, a higher total pressure ratio across the rotor (see Figure 7) results in an increased local choking mass flow rate. This ensures that a reasonable level of mass can still be swallowed, despite a reduction in the effective throat area due to blockage in the diffuser. As the total pressure is rapidly dissipated downstream, there is no global adverse pressure gradient, as discussed in the next section.

A breakdown of the energy conversion mechanisms

Once the kinetic energy (KE) has been absorbed by the flow through the rotor, it must eventually dissipate into internal energy u within the diffuser and vaneless space. In some configurations of the machine, such as the ultra-fast gas preheating zone before the reaction starts, the vaneless space can be removed, and the majority of the energy dissipation occurs every 2–4 stages. In contrast, for a gas mixture undergoing reactions, it is more optimal to dissipate the mechanical energy at every stage to repeatedly balance and compensate for the drop in temperature caused by the reaction.

A well-designed energy transformation system aims to convert kinetic energy into internal energy over the minimum possible time scale, typically less than 0.5 ms. This is achieved controllably through both an isentropic pressure rise and through entropic dissipation mechanisms equivalent to aerodynamic losses. By increasing the entropic rise ds, the increase in static temperature dT can be maximised within a minimum time scale while removing any minor unwanted gains in static pressure dp. This principle can be expressed by the Gibbs equation $c_p dT = (1/\rho)dp + Tds$. Exploiting aerodynamic losses in this way is in stark contrast to conventional energy-imparting machines.

Figure 9 illustrates an approximate breakdown of the energy conversion mechanisms at the two flow regimes shown earlier in Figure 8 above. The energy conversion breakdown is determined by integrating the viscous dissipation rate $\dot{S}_{\mathrm{vsic}}$ (Denton, 1993) zonally for five regions in the diffuser-vaneless-space system listed in Figure 9 (*Left*). A more detailed description of the "loss audit" methodology can be found in Moore and Moore (1983) and Pullan et al. (2004). The sixth energy conversion mechanism accounted for is the isentropic pressure rise $\zeta_{\mathrm{isentropic}}$ (see Equation 2), which is calculated using isentropic relations between two known states. Other conversion sources, such as tip leakage, are neglected.

$$\zeta_{\mathrm{isentropic}} = \frac{c_v \Delta u_{\mathrm{isentropic}}}{h_{0,\mathrm{rotor}} - h_{\mathrm{rotor}}} \tag{2}$$

Equation 3 is used to calculate the entropic portion $\zeta_{\mathrm{entropic}}$ of the total energy conversion coefficient $\zeta = \zeta_{\mathrm{isentropic}} + \zeta_{\mathrm{entropic}}$. It is noted that the subscript $(_{\mathrm{rotor}})$ refers to the rotor exit station.

$$\zeta_{\mathrm{entropic}} = \frac{\dot{S}_{\mathrm{vsic}}}{\dot{m}(h_{0,\mathrm{rotor}} - h_{\mathrm{rotor}})}, \quad \text{where} \quad \dot{S}_{\mathrm{visc}} = \iiint_v \tau_{ij} \frac{dV_i}{dx_j} dv \tag{3}$$

Figure 9 illustrates the differences in the energy conversion mechanisms for two feeds at different reaction states. In the subsonic case (i.e., high-temperature decomposed ammonia), isentropic energy conversion dominates since the flow in the blade passage is more attached. For the supersonic case, there is still a contribution from isentropic pressure rise, but the contribution is smaller, and the mechanism is dominated by pressure rise due to the isentropic component of the shockwave action rather than diffusion induced by changes in the effective flow area.

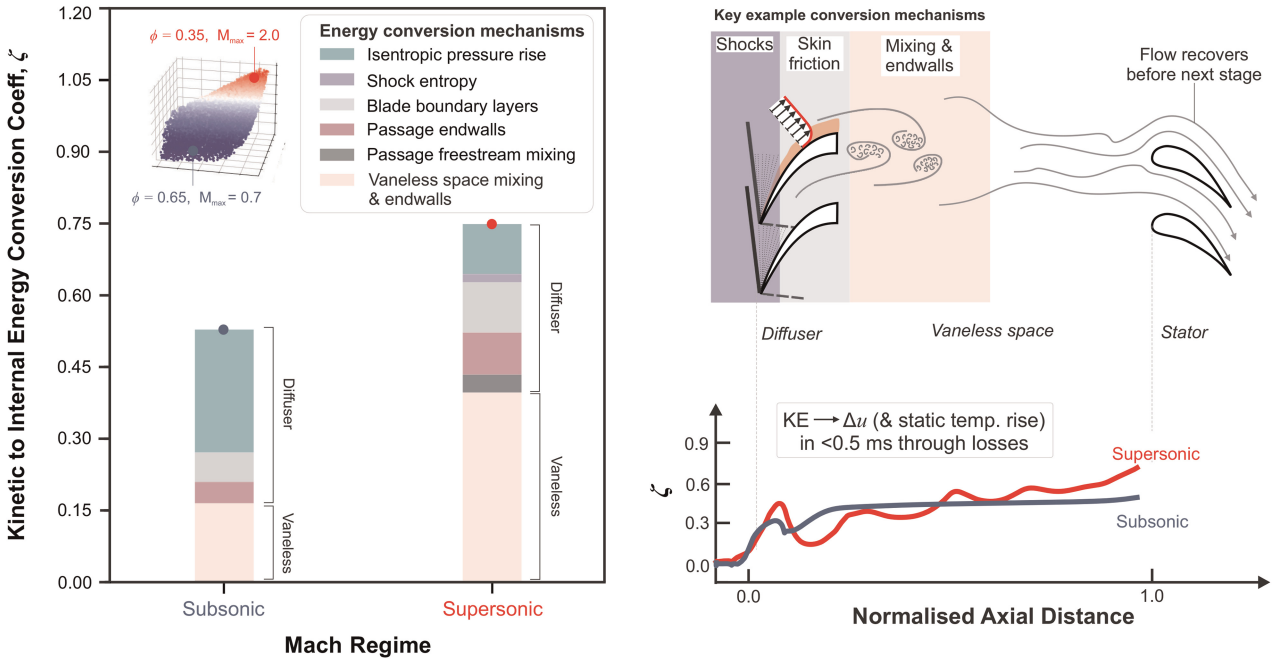


Figure 9. A breakdown of kinetic energy (KE) to internal energy ($\Delta u = c_v \Delta T$) conversion mechanisms including both entropic dissipation and isentropic pressure rise contributions. The breakdown for a single stage is compared for two extreme Mach regimes corresponding to thermophysical properties encountered for two different reactions.

In both cases, the energy conversion due to trailing edge (TE) mixing is large since the TE is thick (42% of the blade pitch), blunt and square. In the supersonic case, this energy conversion source is well over 50% of the total energy conversion. This is significantly higher than the subsonic case as TE mixing loss scales with Mach number Denton and Xu (1990). In the supersonic case, the entropy generated through the shock is relatively small, meaning that the *direct* impact of the shock is minor. However, there is a substantial *indirect* impact from the shockwave due to shock-induced boundary layer separation. This significantly enhances boundary layer dissipation as well as downstream mixing.

This analysis shows that, despite variations in the incoming dynamic head over a wide range of operating points and feeds, almost 60% of the rotor exit kinetic energy can be dissipated into internal energy over a short distance (<50 mm). In different regimes, the flow self-adjusts the contributions from each conversion mechanism to match the downstream boundary condition. Future work will investigate the energy conversion mechanisms in more detail.

Summary and conclusions

This study has demonstrated the feasibility of using an electric-motor-driven turbomachine to mitigate carbon emissions and improve the operability, controllability, and reaction performance of four crucial energy-intensive endothermic reaction processes. The numerical investigation presented in this paper has revealed several key design characteristics of the concept summarised in Figure 10, which are as follows:

1. It has been shown that a universal stage design philosophy can be used for a wide range of feeds and reaction states. The specific energy input remains relatively high over a wide range of conditions. Since the collection of molecular weights covered is extensive, many other relevant chemical and non-chemical processes should fall within this range.
2. The robustness of the turbomachine has been successfully demonstrated over a wide range of operating and Mach number regimes by inducing operating point excursions through fluid property perturbations. The stage loading only dropped by 15% over a broad set of flow conditions, indicating the flexibility of the RDR. Additionally, the turbomachine showed a weak dependence on the Reynolds number, with only a 9% variation in the work coefficient over two orders of magnitude of the Reynolds number.
3. Without a requirement for pressure gain, the design restrictions of efficiency and system instability are no longer relevant. Therefore, a high level of incidence and separation can be endured, facilitating a high level of

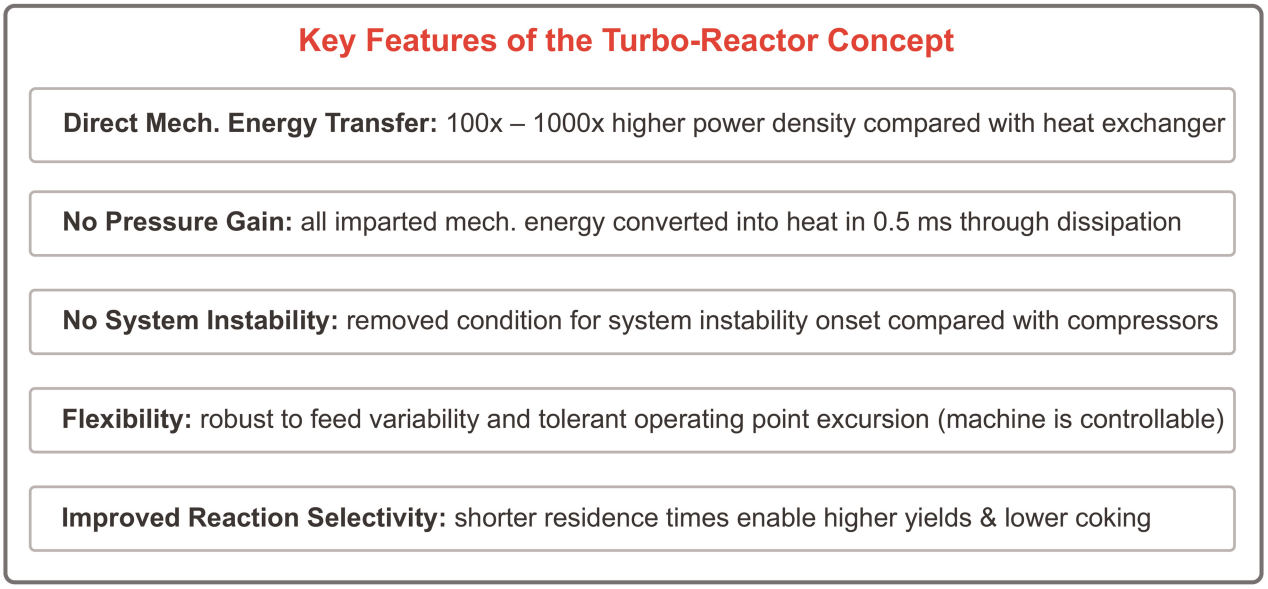


Figure 10. Schematic block diagram summarising the key features of the new RDR concept.

energy input and swallowing capacity over the operating range. The limits of the RDR are set only by choking (which is a restriction on the throughflow capacity of the machine) and not by system instability. Of course, local flow instability within certain blade rows may still occur.

4. Over a broad range of gas velocities and Mach numbers—corresponding to different feeds and reaction states—this work has shown that kinetic energy can be dissipated into internal energy (and static temperature rise) over an extremely short time scale of 0.5 ms, enabling a compact machine for ultra-fast gas heating. The processing mechanisms self-adjust in different flow regimes.

In summary, a turbomachine is a robust and efficient means of heating a wide variety of (endothermically) *reacting* gas mixtures to high temperatures. This analysis covers a broad spectrum of fluid properties, therefore suggesting that a similar elemental stage design could also be used to heat other *non-reacting* gas mixtures, enabling decarbonisation of energy-intensive industries, such as steel, cement, glass, and aluminium, while simultaneously reducing the volume 50 to 1,000 times relative to conventional surface heat exchangers.

Nomenclature

Symbols

b	true blade chord, m
$c = \sqrt{\gamma RT}$	speed of sound, m s^{-1}
$C_{\psi_{TS}} = \frac{p_{exit} - p_{0,in}}{0.5\rho U^2}$	total-to-static pressure rise coefficient
c_p	isobaric heat capacity, J kg^{-1} K^{-1}
c_v	isochoric heat capacity, J kg^{-1} K^{-1}
h_0	stagnation enthalpy, J kg^{-1}
$\Delta_{rxn}H^0_{298}$	enthalpy of reaction, J kg^{-1}
$\dot{m}^*$	non-dimensional choking mass flow rate
p	static pressure, Pa
p_0	stagnation pressure, Pa
R	specific gas constant, J kg^{-1} K^{-1}
$Re = \frac{\rho V b}{\mu}$	Reynolds number
$\dot{S}_{vsic}$	viscous dissipation rate, W
T	static temperature, K
Δu	internal energy change, J kg^{-1}
U	mean blade speed, m s^{-1}
v	domain volume, m^3
V	absolute velocity, ms^{-1}

V_x	axial velocity, ms^{-1}
$\gamma = \frac{c_p}{c_v}$	heat capacity ratio
ζ	energy conversion coefficient
μ	dynamic viscosity, Pa s
$\psi = \frac{\Delta h_0}{U^2}$	stage loading/work coefficient
$\phi = \frac{V_x}{U}$	flow coefficient
τ_{ij}	shear stress tensor, Pa

Acronyms

PFR	plug-flow reactor
PS/SS	pressure/suction surface
RDR	RotoDynamic Reactor (turbo-reactor)
SMR	steam methane reforming
SWBLI	shockwave boundary layer interaction

Acknowledgments

The authors thank Coolbrook Oy for supporting this work. The authors are grateful for fruitful discussions with Tuomas Ouni at Coolbrook Oy.

Funding sources

Coolbrook Oy.

Competing interests

Dylan Rubini declares that he has no conflicts of interest. Nikolas Karefyllidis declares that he has no conflicts of interest. Budimir Rosic declares that he has no conflicts of interest. Liping Xu declares that he has no conflicts of interest. Elina Nauha declares that he has no conflicts of interest.

References

Bender M. (2014). An overview of industrial processes for the production of olefins–C4 hydrocarbons. *ChemBioEng Reviews*. 1 (4): 136–147. https://doi.org/10.1002/cben.201400016.

Bushuev V. A. (2016). Bladed Reactor for the Pyrolysis of Hydrocarbons. U.S. Patent Application No. US 9.494,038 B2. Washington, DC: U.S. Patent and Trademark Office.

Coolbrook Oy. (2022). Coolbrook Oy. Helsinki, Finland, accessed Oct. 15, 2022, https://coolbrook.com.

Delikonstantis E., Scapinello M., and Stefanidis G. D. (2019). Process modeling and evaluation of plasma-assisted ethylene production from methane. *Processes*. 7 (2): 68. https://doi.org/10.3390/pr7020068.

Denton J. D. (1983). An improved time-marching method for turbomachinery flow calculation. *ASME Journal of Engineering for Power*. 105 (3): 514–521. https://doi.org/10.1115/1.3227444.

Denton J. D. (1986). The use of a distributed body force to simulate viscous effects in 3D flow calculations. In Vol. 1 Turbomachinery of Turbo Expo: Power for Land, Sea, and Air. p. V001T01A058. https://doi.org/10.1115/86-GT-144.

Denton J. D. (1992). The calculation of three-dimensional viscous flow through multistage turbomachines. *ASME Journal of Turbomachinery*. 114 (1): 18–26. https://doi.org/10.1115/1.2927983.

Denton J. D. (1993). The 1993 IGTI scholar lecture: loss mechanisms in turbomachines. *Journal of Turbomachinery*. 115 (4): 621–656. https://doi.org/10.1115/1.2929299.

Denton J. D., and Xu L. (1990). The trailing edge loss of transonic turbine blades. *Journal of Turbomachinery*. 112 (2): 277–285. https://doi.org/10.1115/1.2927648.

Dixon S. L., and Hall C. (2013). Fluid Mechanics and Thermodynamics of Turbomachinery. Butterworth-Heinemann: Oxford.

Gao C. W., Allen J. W., Green W. H., and West R. H. (2016). Reaction mechanism generator: automatic construction of chemical kinetic mechanisms. *Computer Physics Communications*. 203: 212–225. https://doi.org/10.1016/j.cpc.2016.02.013.

Goodwin D. G., Moffat H. K., Schoegl I., Speth R. L., and Weber B. W. (2022). Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes. https://www.cantera.org. Version 2.6.0.

Greitzer E. M., Tan C. S., and Graf M. B. (2007). Internal flow: concepts and applications.

IEA (2021). Global Energy Review 2021, Technical Report. International Energy Agency (IEA), Paris.

IEA (2022). Hydrogen, Paris, Technical Report. International Energy Agency (IEA), Paris.

Jackson C., Fothergill K., Gray P., Haroon F., Makhloufi C., et al. (2019). Ammonia to Green Hydrogen Project: Feasibility Study. UK: Ecuity.

Karefyllidis N., Rubini D., Rosic B., Xu L., and Purola V.-M. (2023). A novel axial energy-imparting turbomachine for high-enthalpy gas heating: robustness of the aerodynamic design. *Journal of Turbomachinery*. 146(3): 031005. https://doi.org/10.1115/1.4063928.

Kusenberg M., Zayoud A., Roosen M., Thi H. D., Abbas-Abadi M. S., et al. (2022). A comprehensive experimental investigation of plastic waste pyrolysis oil quality and its dependence on the plastic waste composition. *Fuel Processing Technology*. 227: 107090. https://doi.org/10.1016/j.fuproc.2021.107090.

Lu Q., Zhang B., Yang S., and Peng Z. (2022). Life cycle assessment on energy efficiency of hydrogen fuel cell vehicle in china. *Energy*. 257: 124731. https://doi.org/10.1016/j.energy.2022.124731.

Moore J. and Moore J. G. (1983). Entropy production rates from viscous flow calculations: part I — a turbulent boundary layer flow. In Vol. 1 Turbomachinery of Turbo Expo: Power for Land, Sea, and Air. V001T01A032. https://doi.org/10.1115/83-GT-70.

Pullan G., Denton J., and Curtis E. (2004). Improving the performance of a turbine with low aspect ratio stators by Aft-loading. *Journal of Turbomachinery*. 128 (3): 492–499. https://doi.org/10.1115/1.2182000.

Rostrup-Nielsen J., and Christiansen L. J. (2011). Vol. 10. *Concepts in Syngas manufacture*. London: World Scientific.

Rubini D., Karefyllidis N., Xu L., Rosic B., and Johannesdahl H. (2022a). Accelerating the development of a new turbomachinery concept in an environment with limited resources and experimental data: challenges. In Vol. 86120. Turbo Expo: Power for Land, Sea, and Air. American Society of Mechanical Engineers. V10DT36A001.

Rubini D., Karefyllidis N., Xu L., Rosic B., and Johannesdahl H. (2022b). A new robust regenerative turbo-reactor concept for clean hydrocarbon cracking. *Journal of the Global Power and Propulsion Society*. 6: 135–150. https://doi.org/10.33737/jgpps/150550.

Rubini D., Xu L., Rosic B., and Johannesdahl H. (2021). A new turbomachine for clean and sustainable hydrocarbon cracking. *ASME Journal of Engineering for Gas Turbines and Power*. 144 (2): 021024. https://doi.org/10.1115/1.4052784.

Thiel G. P. and Stark A. K. (2021). To decarbonize industry, we must decarbonize heat. *Joule*. 5 (3): 531–550. https://doi.org/10.1016/j.joule.2020.12.007.

Venkataraman K., Wanat E., and Schmidt L. (2003). Steam reforming of methane and water-gas shift in catalytic wall reactors. *AIChE Journal*. 49 (5): 1277–1284. https://doi.org/10.1002/aic.690490518.

Weydahl T., Jamaluddin J., Seljeskog M., and Anantharaman R. (2013). Pursuing the pre-combustion ccs route in oil refineries–the impact on fired heaters. *Applied Energy*. 102: 833–839. https://doi.org/10.1016/j.apenergy.2012.08.044.

Statement of Authorship for joint/multi-authored papers for PGR thesis

To appear at the end of each thesis chapter submitted as an article/paper

The statement shall describe the candidate's and co-authors' independent research contributions in the thesis publications. For each publication there should exist a complete statement that is to be filled out and signed by the candidate and supervisor **(only required where there isn't already a statement of contribution within the paper itself).**

Title of Paper	A Novel Axial Energy-Imparting Turbomachine for High-Enthalpy Gas Heating: Robustness of the Aerodynamic Design
Publication Status	☐Published ☐ Accepted for Publication ☐Submitted for Publication ☐Unpublished and unsubmitted work written in a manuscript style
Publication Details	Karefyllidis, N., Rubini, D., Rosic, B., Xu, L., and Purola, V. (November 16, 2023). "A Novel Axial Energy-Imparting Turbomachine for High-Enthalpy Gas Heating: Robustness of the Aerodynamic Design." ASME. *J. Turbomach*. March 2024; 146(3): 031005. https://doi.org/10.1115/1.4063928

Student Confirmation

Student Name:	Dylan Rubini		
Contribution to the Paper	Did optimization, methodology modelling, tool development and analysis for chemistry-related section of the paper. Write entire chemistry-related section of paper. Shared majority of analysis, postprocessing and writing workload for the entire paper.		
Signature		Date	16/08/2024

Supervisor Confirmation

By signing the Statement of Authorship, you are certifying that the candidate made a substantial contribution to the publication, and that the description described above is accurate.

Supervisor name and title: Prof. Budimir Rosic		
Supervisor comments I confirm that the candidate's statement is an accurate description of his contribution to this publication.		
Signature 	Date	19/08/2024

This completed form should be included in the thesis, at the end of the relevant chapter.

4.3 Perspective

This chapter has evaluated two features of the robustness of the turbo-reactor concept. This robustness is rooted in the fact that there is no global increase in total-to-static pressure ($p_{\text{exit}}/p_{0,\text{inlet}} \leq 1$), so the onset condition for violent system instability is avoided ($\partial C_{\Psi_{\text{TS}}}/\partial\phi < 0$), and isentropic efficiency constraints are alleviated. Since the mechanical energy imparted within each stage is always closely coupled with rapid dissipation downstream, the available energy that could feed system instabilities is removed. Instabilities can grow locally and propagate circumferentially, but should not be catastrophic. As a result of this robustness, the risks associated with the aerodynamic design are mitigated.

Section 4.1 demonstrated that the elemental stage design is tolerant to interstage disturbances, despite the vaneless space distance being shrunk to one diffuser chord length. This enables the design of a reaction-efficient temperature profile by redistributing the enthalpy input stages. Enhanced reaction efficiency is a key feature aiding the turbo-reactor's cost competitiveness with conventional technology.

Subsequently, Sec. 4.2 demonstrated that the universal elemental stage design was robust over a wide range of feedstocks, chemical reaction processes and reaction states. Despite constructing an artificial "worst case scenario" in which the blade design and boundary conditions are fixed, the machine operated stably with only a 15% variation in enthalpy input observed over a wide range of thermochemical conditions (*e.g.,* $0.5 \leq M \leq 2.0$).

In this chapter, the chemistry is completely decoupled from the aerodynamics. This is a powerful lower-order approach to estimate streamwise variations in aerothermal performance (*e.g.,* stage loading coefficient) without the large computational overhead of reacting flow simulations. However, more accurate aerochemical modelling is necessary at the design level for an efficient reaction to be achieved in practice.

Chapter 5

The Design Ecosystem: New Requirements

5.1 Accelerating the Turbo-Reactor Development

This chapter provides a bridge between chapters focused on aerodynamics and chapters related to tool development to capture aerochemical interactions.

It will be underlined that the flowfield is uniquely complex, featuring large separations, low Reynolds number transitional flow, high freestream turbulence, and a wide range of Mach numbers from low subsonic to high supersonic. This unique level of complexity creates two challenges. First, many databases of experimental correlations may not apply to this novel turbomachine. Second, standard Reynolds-averaged Navier-Stokes (RANS) turbulence treatments struggle with these types of flows. On top of this, chemical reactions are strongly interacting with the complex aerothermodynamics. At the end of this chapter, a very simple method is introduced to capture interactions between the reaction and the energy equation. However, this may not be accurate enough for all stages of the design cycle.

Pages 113–129:

Accelerating the Development of a New Turbomachinery Concept in an Environment with Limited Resources and Experimental Data: Challenges

(*ASME Turbo Expo* 2022)

Proceedings of ASME Turbo Expo 2022
Turbomachinery Technical Conference and Exposition
GT2022
June 13-17, 2022, Rotterdam, The Netherlands

GT2022-80698

ACCELERATING THE DEVELOPMENT OF A NEW TURBOMACHINERY CONCEPT IN AN ENVIRONMENT WITH LIMITED RESOURCES AND EXPERIMENTAL DATA: CHALLENGES AND METHODOLOGY

Dylan Rubini[1*], Nikolas Karefyllidis[1*], Liping Xu[2], Budimir Rosic[1], Harri Johannesdahl[3]

[1]Oxford Thermofluids Institute, University of Oxford, Oxford, OX1 3PJ, UK
[2]Whittle Laboratory, University of Cambridge, Cambridge, CB3 0DY, UK
[3]Coolbrook Oy, Pieni Roobertinkatu 9, 00130 Helsinki, Finland

ABSTRACT

This paper presents the requirements for a tailored design ecosystem to accelerate the development of a new turbo-reactor for hydrocarbon cracking. The objective of this novel turbo-reactor concept is to eliminate carbon emissions from the radiant section of the plant. However, the time constraint for this energy- and carbon-intensive cracking industry to fully decarbonize is severely restricted. Therefore, the development of this new concept must be accelerated to a market-ready product within a limited time frame. However, the novelty of the design requirements and the unique level of complexity of the flow physics compared with conventional turbomachines result in inadequate experimental data, with insufficient time and resources to develop a new measurement database. This shifts the design system towards computer-aided engineering platforms; however for the turbo-reactor, the uncertainty level of these empirically-tuned low-order numerical tools (e.g., RANS and throughflow solvers) is elevated for three reasons. First, blockage, deviation and loss correlations integrated within throughflow solvers are unsuitable to capture the flow physics in the turbo-reactor. Second, RANS turbulence closure models need to be re-calibrated for the complex, low Reynolds number and transonic flow physics in the RDR. Finally, due to the chemical reactions, the working fluid properties vary significantly along the gas path and strongly influence the aerothermodynamics. A new tailored design ecosystem is proposed to address these challenges. The relatively low Reynolds number regime throughout the turbo-reactor enables upfront and routine eddy-resolving simulations to be performed. This allows the new design system to feature a feedback mechanism in which high-fidelity CFD data is used to calibrate lower-fidelity—but less CPU intensive—design tools and empirical correlations. it is demonstrated that the design is robust, enabling a desensitization to modeling errors and uncertainties. Crucially, the new toolchain has been customized to provide the aerodynamic designer with immediate insight into how stage-design modifications influence the reaction product yield and coke deposition rate distribution. This paper demonstrates the capabilities of an efficient aerodynamic-chemistry coupling methodology that has been implemented within an in-house CFD solver, as well as highlighting the potential for reduced-dimensionality kinetic modeling to accelerate optimization of the reaction profile.

NOMENCLATURE

Symbols

W_m	molecular weight of species m
Y_m	mass fraction of species m, %
$\psi = \frac{\Delta h_0}{U^2}$	work/loading coefficient
ϕ	reaction progress, %

Acronyms

CAE	computer-aided engineering
RANS	Reynolds-averaged Navier–Stokes

*Address all correspondence to this author: dylan.rubini@eng.ox.ac.uk

Copyright © 2022 by ASME

LE	leading edge
LES	large eddy simulation
PV	passage vortex
RDR	Roto-Dynamic Reactor
SWBLI	shockwave boundary layer interaction
TLV	tip leakage vortex

1 INTRODUCTION

With the impending threat of irreversible climate change, legislation, regulation, and carbon pricing will inevitably place rigid constraints on the time-scale before which energy-intensive industries, such as the chemical and petrochemical industries, must fully decarbonize. For example in Europe, the "Fit for 55" legislative package published in 2021 proposes a stricter EU Emissions Trading System and carbon border adjustment mechanism [1]. Therefore, this urgent demand to curtail carbon emissions has promoted a race for innovation in these industries.

A new turbo-reactor technology—named the Roto-Dynamic Reactor (RDR)—is introduced as a means of potentially fully decarbonizing the highly energy- and carbon-intensive, high-temperature hydrocarbon cracking process. More specifically, the RDR currently targets a steam cracking process for a range of feedstocks including naphtha, ethane, propane, and biomass and synthetically-derived feeds. Figure 1 shows this novel turbomachinery concept, highlighting how it attempts to eliminate the fundamental limitations of the conventional process. This is achieved by replacing the (indirect) surface heat exchange energy transfer mechanism in the radiant section with *direct* mechanical energy transfer using an electric-motor-driven turbomachine, therefore, directly (and renewably) electrifying the process.

Other low-carbon strategies such as using an electric heater furnace or burning hydrogen enable a reduction in the carbon footprint of steam cracking; however, they provide limited additional benefits. The new turbo-reactor provides innovation in all aspects. Not only does it enable the potential for complete decarbonization using an electric motor powered by renewable electricity, the turbo-reactor could also enable a 20% increase in reaction product yield [2, 3], 30% reduction in energy consumption, 500-fold lower reactor volume, and potentially a dramatic reduction in coke formation and deposition. The knock-on effect of these features is a significant reduction in capital, operating and maintenance costs. Potentially, this could increase the profit of the cracking process by 60%.

In order to out-pace the fierce competition from global oil companies developing electric heating furnaces [4, 5] and meet the pressing demands of decarbonization, the alternative turbo-reactor technology must be accelerated and brought to market within a few years. The challenge is that there is limited experimental data available for this (or a similar) machine, and there is insufficient time and physical resources to develop an entirely new measurement database. This drives the design process towards computer-aided engineering (CAE) systems.

The primary objective of this paper is to highlight the challenges encountered during design of the turbo-reactor, and the new tools and methodologies necessary in a landscape with limited experimental data. This study presents the unique level of complexity resulting from the novel design requirements, ultra-high power density, and widely varying fluid properties (due to the chemical reaction). Therefore, the available empirical correlations and cascade measurement databases, which are in-built into CAE tools, are unsuitable for this machine.

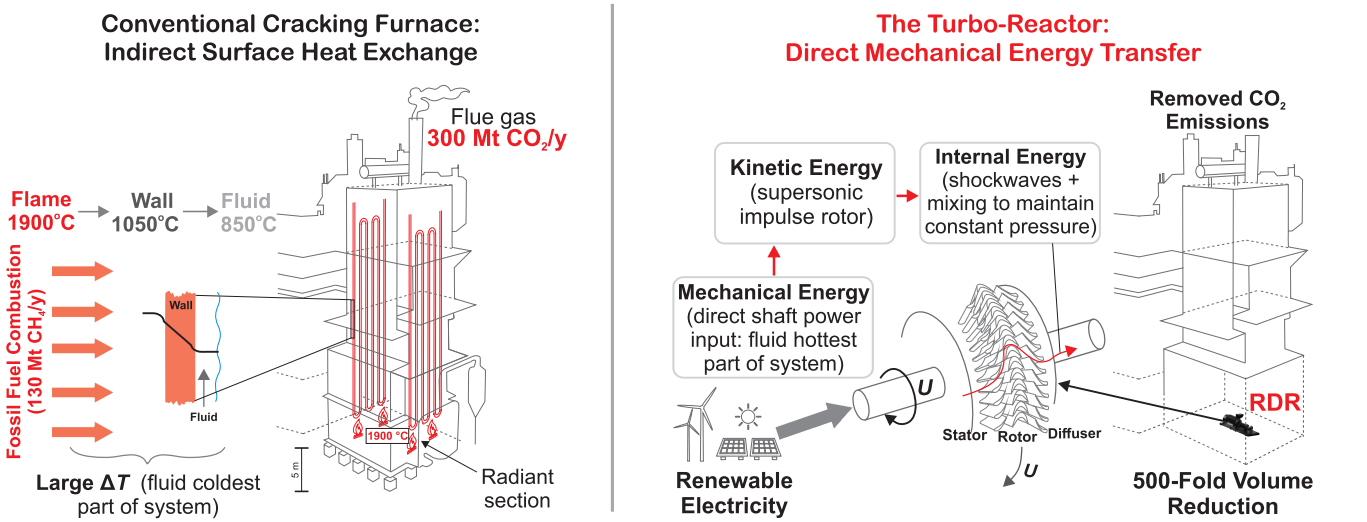


FIGURE 1: A schematic comparison of the energy transfer mechanism within a conventional ethylene production plant (left) and a plant featuring the new turbo-reactor technology (right)

Copyright © 2022 by ASME

However, it is demonstrated that the relatively low Reynolds number regime (50,000 – 200,000) along the gas path enables routine and upfront large eddy simulations (LES) to be deployed. This high-fidelity data, generated by LES, can be used to substitute the experimentally-derived correlations that are traditionally applied within low-order design tools (*e.g.,* RANS and throughflow solvers). This enables fast, automated design and optimization by exploiting the new design ecosystem illustrated in Fig. 2. Vitally, the robustness of the design desensitizes the flow to modeling errors and uncertainties; therefore, this reduces the risk associated with the dependency on numerical tools.

A key feature of this new design system is a collection of reduced-dimensionality kinetic modeling techniques, of varying fidelity, for aerodynamic-chemistry coupling. Crucially, this ensures that (1) the change in fluid properties due to the reaction are accounted for, and (2) the designer can rapidly tailor the aerodynamics to maximize the product yields whilst minimizing coking (*i.e.,* a solid carbon deposit on the surfaces of the RDR).

Since the objective of this paper is focused on the design *challenges*, *requirements* and *methodologies* for a new tailored design ecosystem, the numerical results presented are mainly indicative and illustrative. A more detailed presentation of the numerical results is beyond the scope of this paper.

2 THE TURBO-REACTOR CONCEPT AND THE UNIQUELY COMPLEX FLOW PHYSICS

An ideal environment for efficient, controllable and high-yield cracking requires three vital conditions to be met.

- A high reaction temperature must be achieved whilst reducing non-homogeneity in the temperature field.
- The high temperature level must be attained within a short time and length scale (*i.e.,* short residence time).
- The optimum static pressure level must be maintained along the gas path to minimize low-quality secondary reactions.

For the RDR, it can be demonstrated that these thermokinetic parameters can be tailored with a substantially higher level of control than for conventional radiant coils [6, 7]. Most importantly, the mean reaction temperature can be raised controllably up to 250 °C higher than that for conventional tubes if necessary, and the residence time can be reduced by up to an order of magnitude. Therefore, the turbo-reactor provides a mechanism for a comparatively higher quality, more selective and more efficient hydrocarbon cracking process.

Based on these design requirements, the elemental stage consists of a novel three-blade-row repeating stage design including stator, rotor and diffuser blade rows, as shown in Fig. 3. In order to achieve a multistage work input, the elemental stage (Fig. 3(*a*)) can be configured in several ways. Two potential configurations include a regenerative architecture—as shown in Fig. 3(*c*)—and an axial-type machine architecture, as shown in Fig. 3(*b*).

The proceeding three sections emphasize the inherent level of complexity of the flow physics manifesting from the stage design and cracking requirements. This complexity can be partially attributed to the relatively low Reynolds number transitional regime (the critical Reynolds number being ≈350, 000 [8]) influencing the flow-interaction mechanisms. These fluid dynamic processes pose a unique challenge for traditional numerical tools and prevent the applicability of available empirical correlations. For a more detailed description of the working principles of the elemental stage, the reader is referred to Rubini *et al.* [6, 7]

Stator and Rotor The stator blades (inlet guide vanes) are designed to accelerate, precondition and pre-swirl the flow. Following this, an ultra-high level of enthalpy is imparted within a short time scale (≈50 µs) using a transonic *ultra*-high loading rotor blade-row ($\psi = 3.0 - 4.5$). As a result, the supersonic Mach regime in the turbo-reactor (which is up to M = 2.5 – 3.0) is higher than that experienced in most conventional turbomachines (which is often below Mach 2.0).

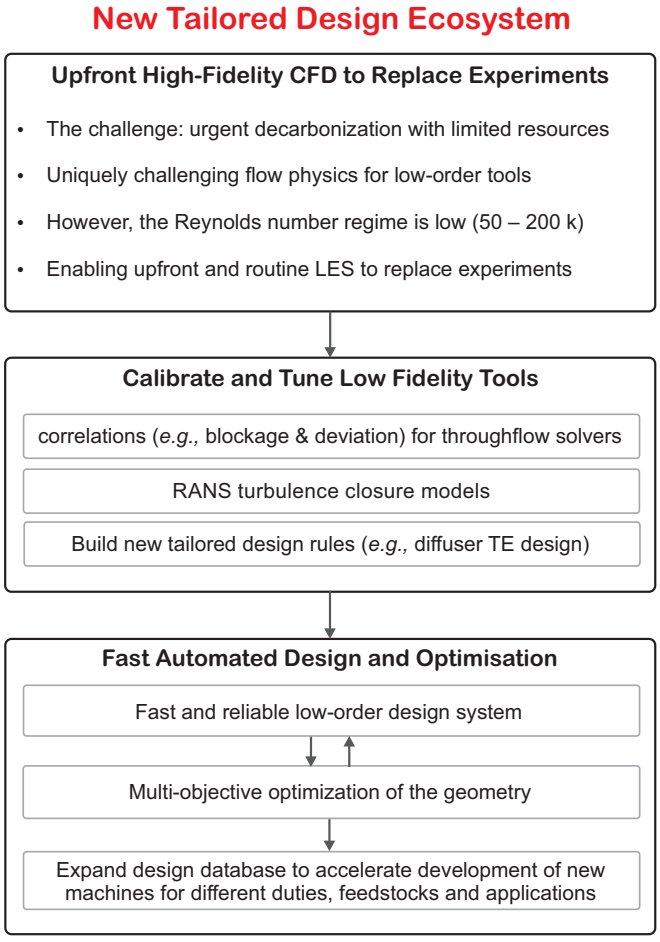


FIGURE 2: **A block diagram of the tailored design ecosystem**

Copyright © 2022 by ASME

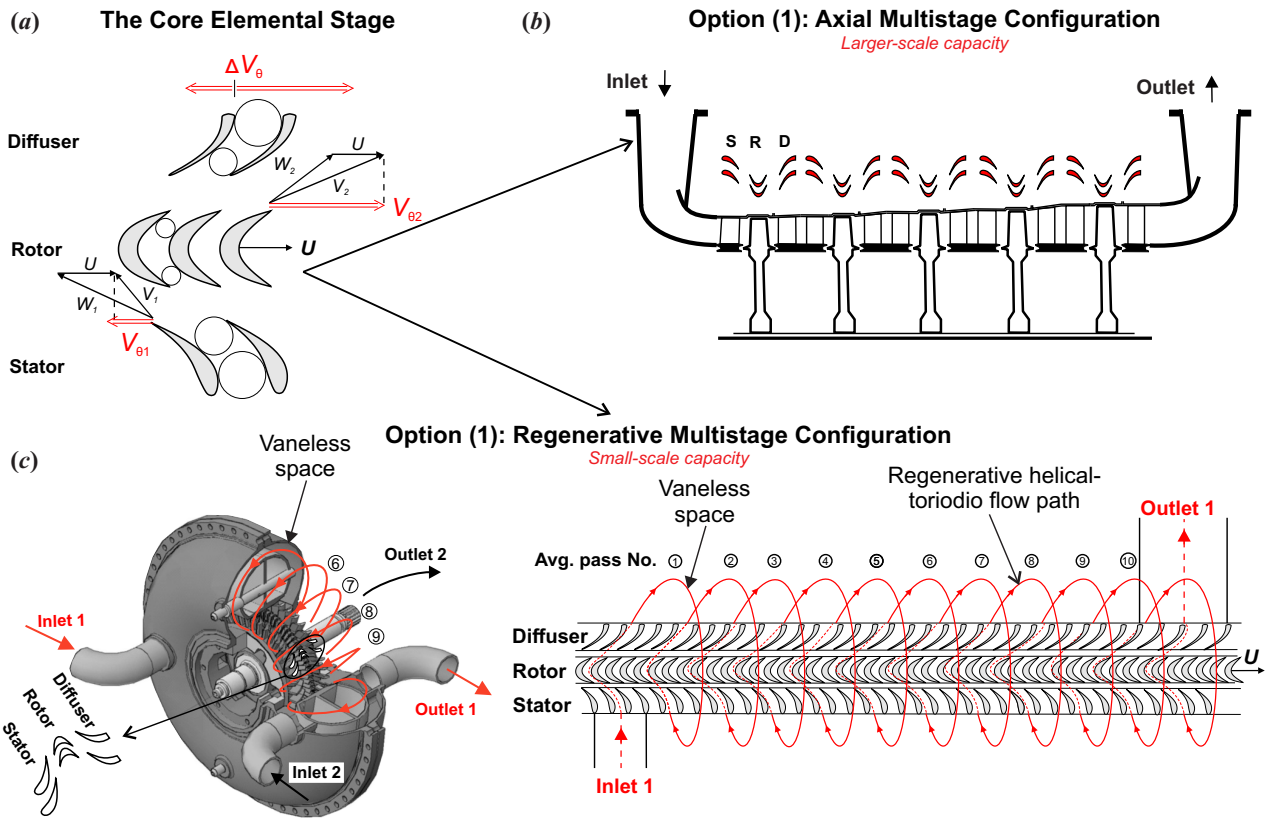


FIGURE 3: **Schematic for (*a*) the elemental stage (*b*) an axial multistage configuration (*c*) a regenerative multistage configuration**

At these supersonic throughflow speeds (see Fig. 4), density changes and gas dynamic effects are appreciable. Figure 5 indicates a strong shockwave at the rotor leading edge (LE). The shock topology can vary depending on the incidence and machine architecture. The shock system strongly influences the flowfield and boundary layer both in the rotor passage, in addition to perturbing the stator boundary layer (see Fig. 5(*a*)).

Figure 5 illustrates that the flow physics in the rotor is dominated by shockwave boundary layer interactions (SWBLI) that control the level and intensity of time-dependent flow separation in the passage. High surface curvature (and flow turning) and negative incidence on the pressure surface amplifies the sensitivity of the boundary layer to detachment (see Fig. 5). Furthermore, as a direct consequence of the low Reynolds number regime (50,000 – 200,000), the thick boundary layers are susceptible to separation and the associated bubble lengths are much longer than those encountered in typical aeroengines [9]. As a result, large-scale unsteady flow separation sets both the effective flow-area—therefore impacting the blockage and hence flow capacity—as well as fixing the level of rotor flow turning.

The transitional nature of the low Reynolds number flow, coupled with the high Mach regime results in a flowfield sensitive to inlet Mach number and incidence perturbations. In the rotor, these effects can (1) change the location of the shedding shear layers (Fig. 5(*a*)) in the passage (2) change the shock position and/or its strength and (3) introduce additional instantaneous passage shockwave. In this highly non-linear and sensitive flow regime, "real geometry" features such as the finite thickness of the LE can drastically change the shock position and strength, transition location, and incidence. Therefore, these real-geometry features can dramatically impact the capacity and loading.

Figure 5(*b*) also illustrates that a combination of high flow turning (≈130°), a large blade-surface pressure differential, and low aspect ratio blading, lead to a highly three-dimensional flowfield in the rotor. The resulting vortical structures take the form of intense secondary flows penetrating deep up the span and strong over-tip leakage flow. As a result, the flow turning capability is severely reduced due to large flow deviation at the blade TE.

Diffuser The transonic diffuser is multi-functional, but its primary function is to re-turn the flow (almost) axially and transform kinetic energy into internal energy for the reaction whilst also ensuring the optimum static pressure level for cracking is maintained. This is achieved through an intense turbulent mixing process, which is encouraged by shock-induced flow separation

Copyright © 2022 by ASME

(see Fig. 5(*a*)). The shock-induced temperature rise provides a crucial mechanism enabling energy transformation within such a short length and time scale. For a heavy hydrocarbon-steam mixture, over ≈100 °C static temperature rise can be achieved in ≈150 µs. This contributes towards a reduction in the residence time, and therefore a yield improvement and coking reduction.

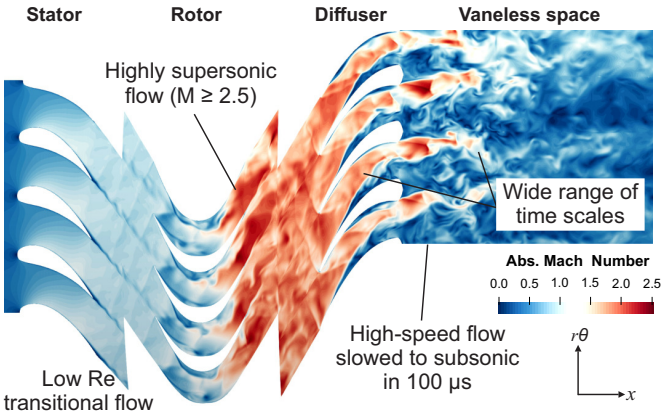


FIGURE 4: **Absolute Mach number contours at midspan illustrating the complexity of the transonic and low Reynolds number transitional flow (LES)**

At the diffuser inlet plane, unsteady vortices shed from the upstream rotor row lead to a highly time-dependent flowfield. Temporal variations in the Mach number and incidence angle result in large excursions in the diffuser shock topology. This controls the time-dependent level separation and blockage in the passage due to SWBLI (see Fig. 5(*a*)).

Vaneless space In the vaneless space connecting stages (see Fig. 3(*b*) & (*c*)), high-energy turbulent mixing is generated using a thick diffuser TE. This dissipation process ensures the optimum cracking pressure is maintained. Furthermore, this enables the different chemical species and non-uniform temperature field to fully mix-out and produce the ideal homogeneous environment for the reaction. There is sufficient space in the vaneless diffuser to complete the mixing process before allowing the flow to recover uniformity—within a few diffuser chord lengths—ahead of the next stage. The relatively low Reynolds number regime often allows the flow to return to a laminar-like state within the vaneless space (or downstream stator).

The preceding sections have underlined how the unique design requirements and operating conditions of the RDR have forged new avenues for highly-loaded, low-reaction transonic stages. The development of this new design paradigm, triggered by the RDR technology, has introduced new complexities into the design and flow physics not present for conventional turbomachines. As a result of this unique fluid dynamics, the uncertainty in the capabilities of standard numerical tools is increased.

3 DESIGNING THE NEW TURBO-REACTOR IN A LANDSCAPE WITH LIMITED EXPERIMENTAL DATA

Traditional turbomachinery design has relied on a wealth of manufacturer experience, know-how, experimental test cases, correlation databases [10, 11], and calibration data to inform the process. However for the new turbo-reactor, the design requirements (high temperatures within short time-scales, and constant pressure), working fluid properties variations due to the reaction, and operating conditions result in a unique elemental stage featuring high-complexity fluid dynamic interactions. This leads to two distinct challenges:

(1) A lack of manufacturer experience and cascade data for preliminary design, detailed design and validation (see Sec. 3).
(2) A high level of uncertainty in the empirical data built into low-order numerical modeling tools (see Sec. 3.1).

Both the rotor and diffuser present uniqueness in their design and associated flow physics. First, the rotor blade-row exploits ultra-high loading (ψ = 3.5 – 4.0) impulsive blading coupled with both supersonic throughflow speeds (up to M = 2.5 – 3.0) within a low stage reaction (*i.e.,* $\Lambda \approx 0$). Second in contrast with conventional blade design, the diffuser is designed to benefit from a controlled level of flow separation and mixing. Therefore, conventional limitations for highly-loaded stages—due to large-scale flow separation in the stationary blade-row—are removed.

Although supersonic impulse blading has not shown widespread use in conventional turbomachines, there are a few instances where the application of this class of blading is ideal. These include, steam turbines [12], Organic Rankine Cycles turbines [13, 14], turbines for driving liquid-fuel propellant pumps within rockets, or helium compression refrigeration pumps [15]. However, studies on these type of rotor blades are commonly for turbines (not work-input devices such as the RDR) and rarely include designs with flow turning greater than 90° [16]. This is substantially less than the turning in the RDR (≈130°). Furthermore, the limited cascade data that is available is for impulse *turbines* [17] where the operating Reynolds number ($\mathrm{Re} \approx 2\times10^6$) is 10 – 50× higher than that in the RDR. Therefore, this cascade data is potentially unsuitable.

The vast majority of modern cascade measurement databases are based on axial compressors/turbines (for aerospace and power generation applications). For axial compressors—even those with stator blade-rows employing flow control techniques (*e.g.,* aspiration) [18]—the loading coefficient is rarely above ≈0.8 [19], and the stage reaction is primarily in the range of 50 – 65% [20]. These parameters are well outside the range that the RDR rotor blades are designed for. The lack of cascade data is even more acute for the multi-functional diffuser. In this blade-row, the requirements are entirely new; therefore, there are limited (or no)

Copyright © 2022 by ASME

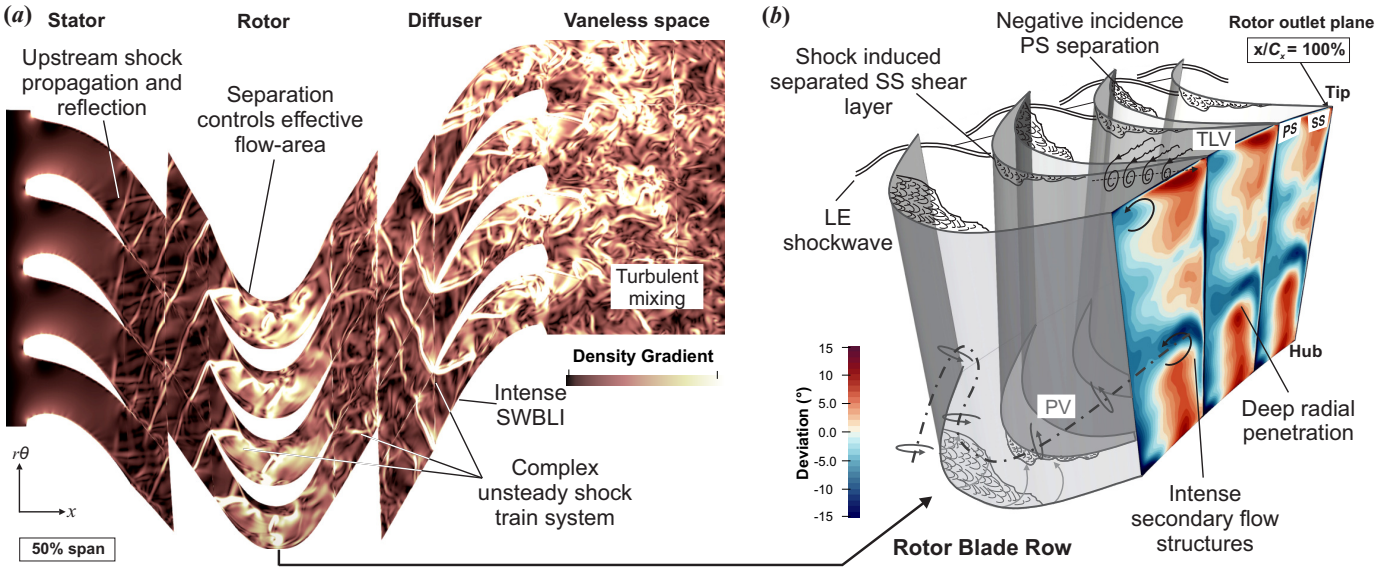


FIGURE 5: **Illustrating the level of complexity of the three-dimensional flow physics (LES) (*a*) across a stage and (*b*) in the rotor**

available studies to be referenced for this design.

3.1 Challenges for Low-order Blade Design Systems

It has been highlighted that the availability of substantial cascade measurement data is limited throughout the design cycle. Therefore, the turbomachinery designer relies predominantly on CAE tools throughout the design process. However, the complexity of the flow physics in the turbo-reactor means that low-order numerical tools may not be appropriate without judicious correction. These tools must be rectified to establish accurate prediction of the three most fundamental design parameters for the RDR: flow turning, throughflow capacity, and mixing intensity.

The limitations of traditional design systems for numerically modeling the unique physics in the RDR are two-fold. First, throughflow solvers struggle with transonic flow and contain in-built blockage, deviation and loss correlations that are unsuitable for this machine. Second, Reynolds-averaged Navier–Stokes (RANS) turbulence closure models are inadequate to produce sufficiently accurate predictions of the complex flows generated by the turbo-reactor design.

2D Throughflow Solvers For 2D low-order throughflow solvers, the critical challenge is reliably capturing the effects of the dominating flow physics in the RDR. These include:

- Large-scale flow separation and hence blockage
- Supersonic flow and Mach effects (see Figs. 4 & 5)
- Strong three-dimensionality (*e.g.,* secondary flows, see Fig. 5(*b*)) and hence deviation

Both streamline curvature and matrix inversion techniques either do not or do not accurately predict transonic flow and the associated Mach, incidence, and deviation effects [21, 22]. At this high loading condition, large-scale separation (and hence blockage), turbulent mixing, and strong three-dimensionality deem critical empirical correlations applied within the solver inappropriate. Therefore in analysis mode, there is a high uncertainty level in the work input and flow capacity predictions.

To combat these challenges there are are two avenues. One option is to rely on three-dimensional viscous computational fluid dynamics (CFD) even for early design, and the second option is to build tailored data-driven correlations derived from high-fidelity CFD (as discussed in Sec. 4).

3D RANS The modeling uncertainty of RANS turbulence closure models increases in the harsh environment generated by the turbo-reactor design. The defects of RANS turbulence models are most severe due to the inherently complex flow features. This key challenging physics includes but is not limited to:

- The strong anisotropy and highly three-dimensional elevated strain rates (*e.g.,* due to corner separation in the diffuser and stretching of coherent vortical structures in the vaneless space)
- The low Reynolds number transitional flow regime and regions of non-equilibrium turbulence (these regions can be extended in this low Reynolds number regime)
- The large-scale flow separation and blockage within the blade passages
- The potential relaminarization in the vaneless space
- The elevated freestream turbulence and mixing in the diffuser and vaneless space
- The impact of the finite-rate chemistry on turbulence

Copyright © 2022 by ASME

3.2 The Robustness of the Turbo-Reactor Design: Desensitization to Modeling Uncertainties

Despite the difficulties predicting the flowfield for conventional numerical tools, the risk associated with the heavy reliance on CFD (rather than test rigs) is lower for the RDR compared with, for example, for conventional axial compressors (also a work-input machine).

Conventional turbomachines, are typically very sensitive to errors in incidence and deviation angles at off-design conditions [23, 24]. This intolerance is a consequence of the machine supporting extreme adverse pressure ratios of up to 60:1. Therefore, incidence effects (*e.g.,* separation and blockage) are amplified and compounded across each stage of a multistage compressor (which can have as many as 15 – 22 stages). As a result, violent instability can have damaging consequences to the aerodynamic flowfield, and in the worst case surge can be destructive to the structural integrity of the machine.

In contrast for the RDR, the robustness of the design, intrinsic flow physics, and operating conditions ease the sensitivity to numerical modeling uncertainties for two reasons.

(1) Figure 6 illustrates that the static pressure rise across the machine can be controlled at to a negligible level (as required by the reaction). This is facilitated by an intense turbulent mixing process in the diffuser and vaneless space, enabling a constant static pressure level to be maintained during the ultra-short time-scale (≈100 µs) energy transformation process (see Fig. 4 & 5). Therefore despite the RDR being an energy imparting device, there is no condition for pressure energy to accumulate downstream of the machine nor be stored in a large-volume plenum. Rather, the volumes are distributed in the spaces between the many stages (therefore eliminating the onset condition for system instability).

(2) The stages are coupled with a vaneless mixing/reacting space where large-scale non-uniformities can mix out and dissipate (see Fig. 4 & 5). Figure 7 (left) indicates that the mixing intensity decays rapidly, before flow homogeneity is quickly reset within only a few diffuser C_x downstream into the vaneless space. This fast dissipation rate is a consequence of the low Reynolds number regime, enabling the flow to reset to an almost laminar-like state. Any residual unsteady distortion presented to the stator inlet of the subsequent stage—especially at off-design conditions—is absorbed by the strong favorable pressure gradient within the stator blade-row, as illustrated in Fig. 7 (middle & right). Separating and isolating the stages with a mixing space followed by a highly accelerating stator—with a thick LE (*i.e.,* tolerant to incidence)—prevents the effects of incidence amplifying across the machine (see Fig. 3(*b*)) at off-design (or design) conditions. This allows the flow to recondition and reset between stages despite the distorted and potentially non-uniform upstream environment. It has been demonstrated that the design of this energy imparting machine is sufficiently robust in spite of the simple, uniform and hence low cost stage, annulus and blade design for the RDR.

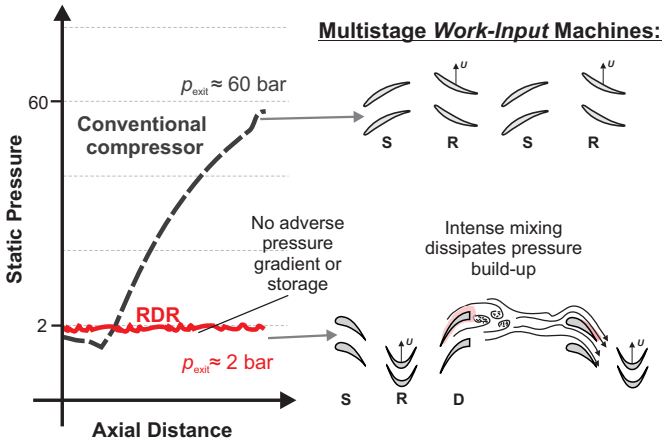


FIGURE 6: **A representative comparison of the pressure accumulation in two multistage work-input machines: (1) the RDR (red) (2) a conventional axial compressor (black)**

For these three reasons, the turbo-reactor is more forgiving to incidence effects and thus less sensitive to the onset of large-scale system instability (*e.g.,* compressor surge). Therefore during the design process for the RDR, it is more acceptable to rely primarily on CFD for detailed design despite the risks of numerical modeling errors and uncertainties.

4 ROUTINE HIGH-FIDELITY LES FOR PRELIMINARY STAGE DESIGN

Having presented the critical challenges in predicting the flowfield, this section highlights the potential to employ routine high-fidelity LES to circumvent expensive (and extensive) experimental measurements. Although reports from Arroyo *et al.* [25] have illustrated successful *one-off* LES coupled into simulations of an entire gas turbine, scale-resolving simulations have yet to become a *routine* tool within a classical industrial multistage turbomachinery design system [26]. In turbomachinery, scale-resolving simulations have been most commonly used to investigate the interactions between the combustor and turbine [27, 28] or for detailed boundary layer investigations in low-pressure turbine (LPT) vanes [29, 30].

For the RDR, the uniqueness of the design cycle is that the Reynolds number is relatively low, in the order of 200,000 at the inlet (before cracking) to 50,000 at the outlet (after cracking). This regime is within a similar range to that experienced in the LPT at high altitude cruise within a small turbo-fan engine [8].

Since the grid-point resolution requirements are a linear function of the Reynolds number ($N_{\text{total}} \propto \text{Re}$ [31]) for wall-modeled LES, it is therefore tractable to resolve sufficient turbulent ki-

Copyright © 2022 by ASME

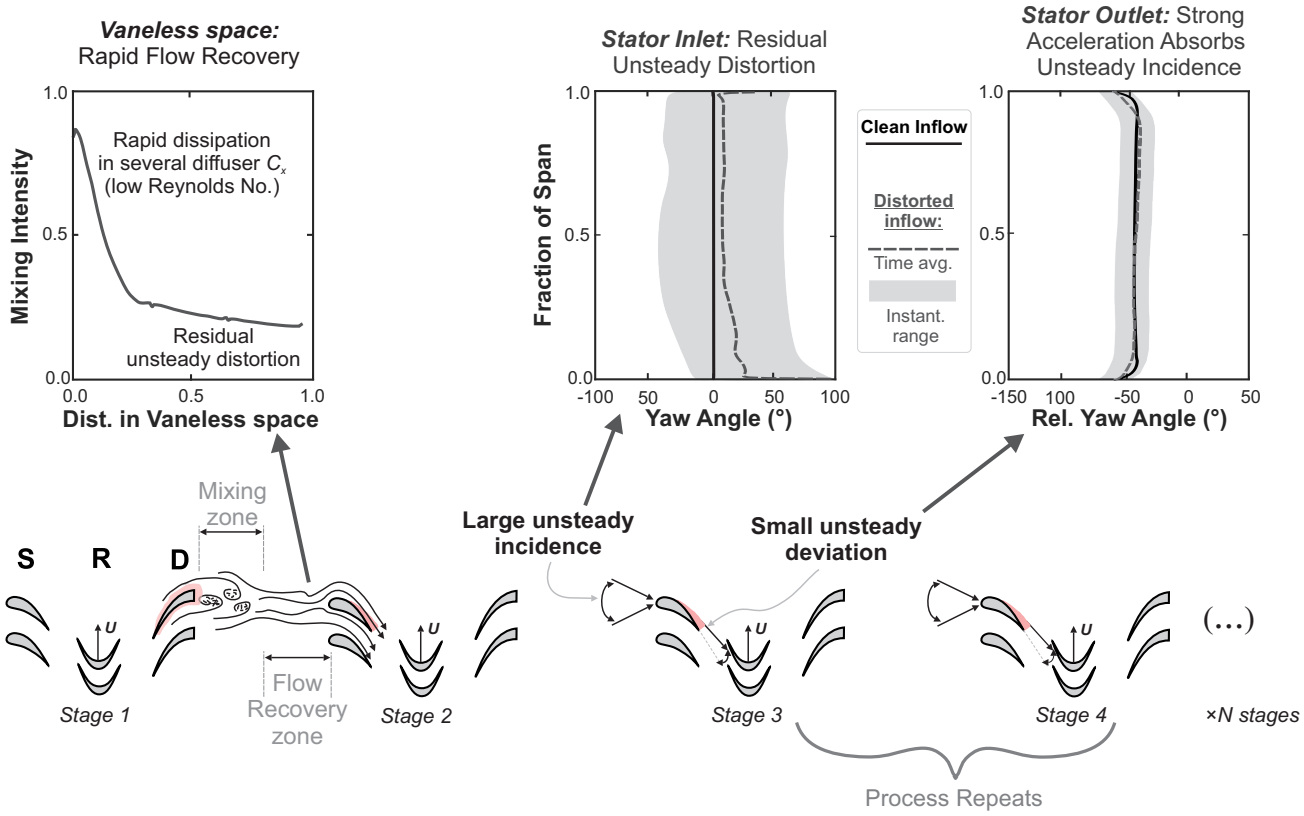


FIGURE 7: **A schematic illustrating the robustness of the design to numerical modeling errors and uncertainties for an energy imparting machine. The intrinsic nature of the design and flow physics mitigates amplification of incidence, deviation and blockage across a multistage machine**

netic energy (*i.e.,* ≈80%) to perform LES upfront and routinely throughout the design cycle. A 1.5 stage domain to be modeled with only a moderate grid density (100 – 200 million) using model high performance computing technology. In this case, the majority of the boundary layer (down to the sublayer) is resolved. It is vital that the eddy-resolving calculations are multi-blade-row (and even multistage) in order to capture the bi-directional blade-row interactions essential for flow physics driving the turbo-reactor.

Supporting (U)RANS calculations with regular high-fidelity CFD provides more confidence in the solution results, enabling the detailed design to rapidly progress from initial concept to market-ready product without substantial time-consuming and costly experimental campaigns. The authors acknowledge that LES cannot replace experimental data nor provide results with as low a level of uncertainty as experiential data; however, due to the pressing demands of decarbonization, it is an attractive alternative that can reduce the design uncertainty (compared with just isolated (U)RANS) without incurring a large computational overhead and/or costs. However, over the next few years, the database of experimental measurements for the RDR will expand and begin to play an increasingly important role.

4.1 Calibrating Lower-Order Numerical models with LES

Throughflow Scale-resolving simulations can be used to tune lower-fidelity, but considerably computationally cheaper, throughflow solver. Therefore, the predictive capability of the low-order tool is upgraded and the design cycle can be accelerated more reliably. The objective is to alleviate the most significant limitation of modern throughflow calculations: modeling the effects of blockage, deviation, loss, and spanwise loss distribution through empirical correlations [32–34]. Eddy-resolving calculations can be used to build new and tailored correlations for these parameters that can then be incorporated into the solver. Furthermore based on the high-fidelity data, it is possible tune the empirical mixing model (within the solver) to account for the intense turbulent mixing process (see Fig. 5(*a*)).

Since the runtime for a throughflow solver can be within a fraction of a second (compared to hours with RANS), a highly-tuned 2D method can be an attractive surrogate solver. There-

Copyright © 2022 by ASME

fore, it is an ideal candidate to accelerate multistage optimization and off-design studies that require expensive CFD simulations for hundreds of different design conditions and/or blade parameters [35]. The throughflow solver can be particularly important stage matching/coupling calculations and for determining the axial length of the inter-stage vaneless space.

Most importantly, once the reduced-dimensionality kinetic model has been incorporated into the throughflow solver (see Sec. 5.1), it will be the ideal candidate to perform multiobjective constrained design optimization to determine the temperature, pressure, residence time and steam injection profiles that. In this case, the objective functions are (1) to maximize the weighted sum of the yields of the primary products (weighted by relative market price/importance) and (2) maximizing the end-of-run time-scale before plant shutdown to remove the thick coke deposit on the walls (approximately every 100 days). This optimization may need to be performed for various feedstocks to account for the feed flexibility of the machine; therefore, many costly reacting flow simulations are required.

RANS Taking this hybrid-fidelity feedback methodology one step further, scale-resolving simulations can be used to tackle the most significant limitation of RANS calculations: turbulence closure and transition modeling. As a means of addressing the turbulence modeling deficiencies highlighted in Sec. 3.1, Fig. 8 compares three pathways currently being investigated aiming to improve the treatment of turbulence whilst maintain a low computational overhead and grid density. The full details of the performance of these extended S-A treatments will be elaborated upon in an upcoming publication. This will highlight the dominant fluid dynamic mechanism that lead to deficiencies in baseline RANS closure modeling. As an example of the first treatment, Fig. 9 illustrates the calibration of the S-A model using LES data. This is achieved by a sparse optimization of the scaling factors pre-multiplying the production, destruction and diffusion source terms in the transport equation relative to the standard default values in for the original model [36].

During this process, the objective function was minimized over a set of 150 RANS calculations. This cost function was defined as the mass-averaged relative error between the RANS and LES solution along a distribution of ≈900 control planes in the streamwise direction. The modifications to the destruction and diffusion scaling factors were minor and within the recommended range [36]. The most significant changes were made to the production term. Currently the amended coefficients are applied globally throughout the domain, indicating the robustness of the model.

Following this model tuning, Fig. 9 indicates that separation and blockage are better captured in the blade rows. Turbulent mixing and dissipation are also more accurately predicted in the diffuser and vaneless space. For the calibrated model, the mass-averaged error in stage loading and flow coefficient (capacity)

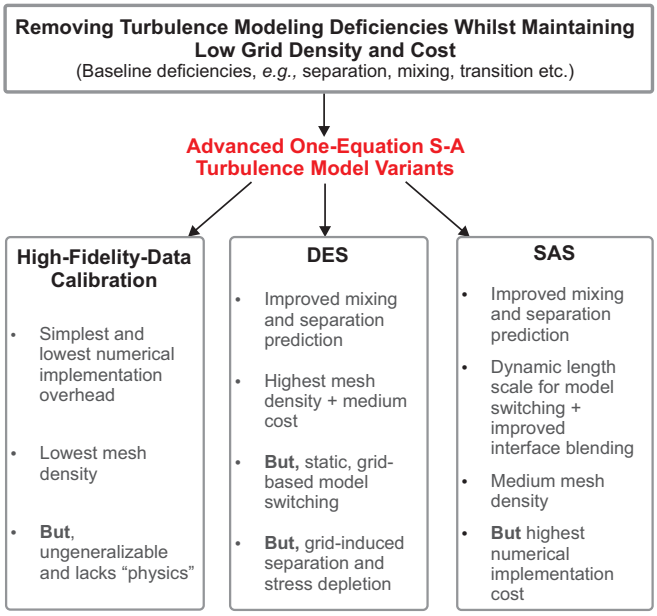


FIGURE 8: **A comparison of advanced S-A turbulence modeling variants being investigated to improve turbulence treatment whilst maintaining a low grid density: (i) high-fidelity-data calibration model (ii) detached eddy simulation (DES) [37] and (iii) scale-adaptive simulation (SAS) [38]**

between RANS and LES has reduced from 3% to 1.5% and 7% to 2.0% respectively.

Although the standard S-A model has been used widely for industrial turbomachinery design, it has been demonstrated that for this novel application there is notable predictive error. The coefficients in the original S-A model were derived and calibrated for lower stage loading airfoils [36] (and canonical flows: simple wakes, free shear layers, mixing layers, well-attached flow, etc.) and are therefore more suitable for the flow physics in traditional and less aggressive turbomachines. For high-lift airfoils, such as those in the RDR, the uncertainty in capturing the onset, length and width of separated regions of the flow can dramatically increase [39]. Furthermore, the applicability of the un-calibrated model is degraded due to several critical factors: (1) the strain rates in the turbo-reactor are significantly elevated and anisotropic (2) the flow is in a transitional regime (3) and there is a strong and concentrated adverse pressure gradient due to intense SWBLI (see Fig. 5(a)), which is then followed by non-equilibrium shock recovery.

For the RDR, the cost function of optimizing yield—the key metric for the quality of the chemical reaction—can be seen as analogous to the conventional objective function of minimizing the efficiency in traditional turbomachinery design. Since the chemical reaction rate is highly a non-linear function of temperature (this will be highlighted in Fig. 12(a) in Sec. 5.1), minor

Copyright © 2022 by ASME

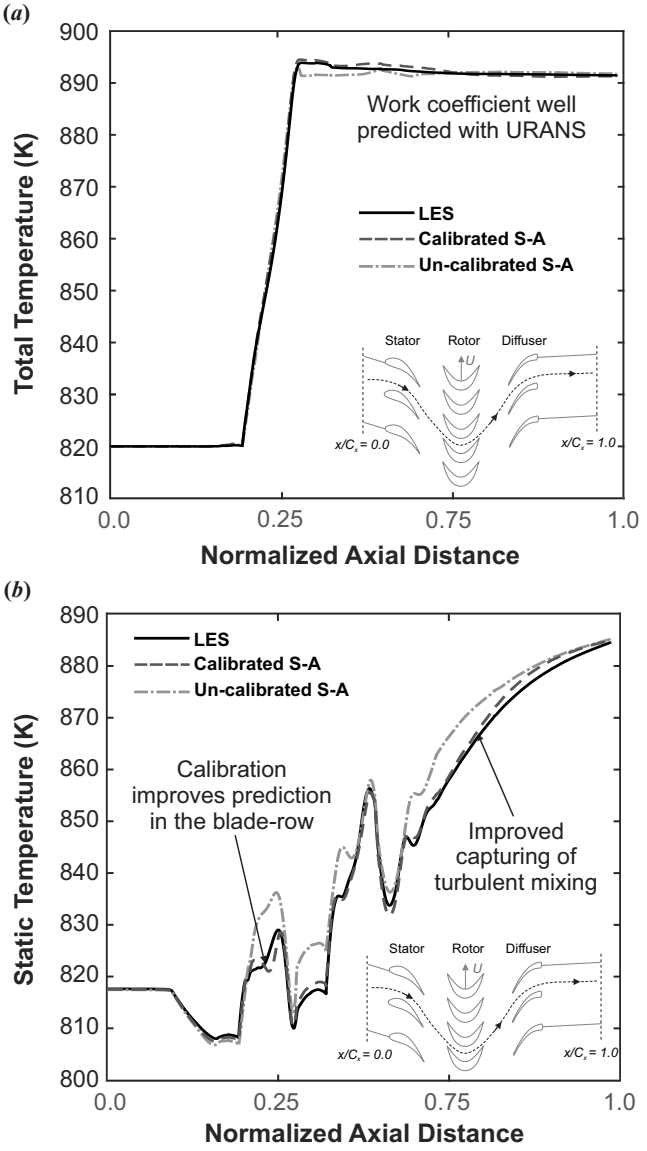


FIGURE 9: **Time- and mass-averaged LES compared with both calibrated and un-calibrated S-A model for (*a*) total temperature and (*b*) static temperature along a single-stage flow path**

uncertainty in the thermokinetic flow parameters can lead to a significant distortion in the yield distribution. Using LES and S-A model calibration, the epistemic uncertainty for turbulence modeling can be systematically reduced, enabling the potential to reduce the energy consumption per tonnes of ethylene produced through careful numerical analysis.

4.2 Building new design rules for the RDR

In the process of parametrically exploring the design space, high-fidelity LES is routinely used to complement URANS calculations. For conventional compressors, extensive databases of design rules are available for optimized airfoil sections [40]. In some instances, these are inappropriate for the RDR. Instead, high-fidelity CFD simulations can be used to build new databases tailored for the turbo-reactor. New design rules have been developed for blade parameters such as aspect ratio, pitch-to-chord ratio (solidity), maximum thickness-to-chord ratio, and diffuser trailing edge (TE) thickness, to name a few.

Large eddy simulations are particularly valuable to build new design rules for the novel multi-function diffuser. In this blade-row, the physics is challenging for low-order tools to predict and the design requirements are entirely new. To address this, eddy-resolving simulations have been used to find the optimum solidity and TE thickness. These blade parameters control the length, location and intensity of the high-energy mixing process (see Fig. 5(*a*)). In conjunction with this, the diffuser TE thickness determines the length-scale of the coherent vortical structures propagating into the vaneless space. Evidently, the objective function used to optimize these parameters are polar to conventional turbomachinery design (*e.g.,* minimize profile less, separation, etc.). Optimizing these parameters without cascade measurements relies on LES calculations for precise prediction of the important mixing process.

5 DEVELOPING A TAILORED DESIGN ECOSYSTEM FOR THE RDR

A core requirement of the new design ecosystem proposed in this paper (see Fig. 2) is a hybrid low- and high-fidelity design toolchain (see Sec. 4). In order to perform routine high-fidelity simulations, it is necessary to address the high computational intensity. This requires highly scalable parallelization for heterogeneous state-of-the-art back-end hardware (*e.g.,* CPUs *and* GPUs) and programming paradigms, as shown in Fig. 10. An embedded domain specific (eDSL) and "active" library (named OP2) facilitate automated parallel code generation for a range of back-end hardware within a single, unified source code [41], and with minimal domain scientist intervention. This parallelization platform is integrated into a new unstructured solver named Utblock. The solver is inspired by a set of established and widely validated structured mesh solvers developed by John Denton [42, 43]. Utblock retains the key features and design philosophy from these structured mesh codes.

Unstructured grid generation facilitates fast, automated and high quality meshing of arbitrarily complex geometries. Coupling the unstructured solver with an automated parallelization platform, enables both accelerated developmental and design turnaround time (aiding in meeting the requirements of urgent decarbonization).

The new design ecosystem and evolving database of tailored correlations enable custom commercial-scale turbo-reactor architectures to be developed automatically and with minimal user intervention. These custom machines must be developed for dif-

Copyright © 2022 by ASME

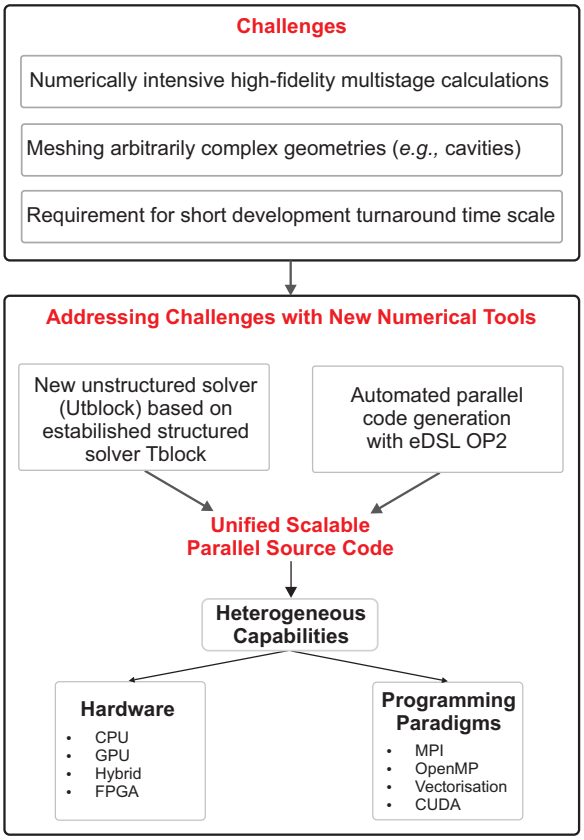


FIGURE 10: **Scalable parallelization system for unstructured mesh applications on heterogeneous back-end hardware**

ferent duties (*i.e.,* mass flow and isentropic work input), varying feedstock composition and new applications.

5.1 Aerodynamic–Chemistry Coupling

An important new tool within the tailored design ecosystem is a mechanism for aerodynamic-chemistry coupling. From an aerothermal perspective, this is important to accurately predict the wide variations in density, temperature, and fluid properties due to the reaction. Therefore, the number of stages, blade design and annulus walls can be optimally configured. On the other side, integrating a low-order reaction model into the aerodynamic simulations will enable multiobjective design optimization of the temperature, residence time, pressure, and steam injection distributions in order to maximize a weighted sum of the primary product yields (weighted based on market price/importance) whilst minimizing the coke (solid carbon deposit) deposition rate (or more appropriately the end-of-run time-scale before decoking requirement necessitate a plant shutdown).

The key interactions and coupling mechanisms between the chemical reaction and aerodynamic flowfield are via *bi-directional* variations in the thermochemical (*e.g.,* gas constant, reaction rates, enthalpy of formation, etc.) and thermophysical properties. For example from one coupling direction, as the gas mixture is cracked and the molecular weight reduces, the gas constant increases by over ≈70%. This large change in the gas constant contributes to a ≈60% reduction in density along the gas path. It is vital that this density variation is accounted for to ensure that (1) the pressure is controlled at the optimum level for cracking and (2) the velocity triangles are uniform across all stages of the turbo-reactor by maintaining a fixed the axial velocity component. From the other coupling direction, static temperature is the primary driver for non-linear variations in the rate constant (and hence species net molar production /destruction rates).

To incorporate the effects of chemistry with a varying level of complexity and fidelity, a three-tier hierarchical approach has been investigated; this is in summarized in Fig. 11. The route taken depends on whether the designer is interested in accurate yield and selectivity estimates or fast aerodynamic design. Therefore, these pathways are being developed simultaneously.

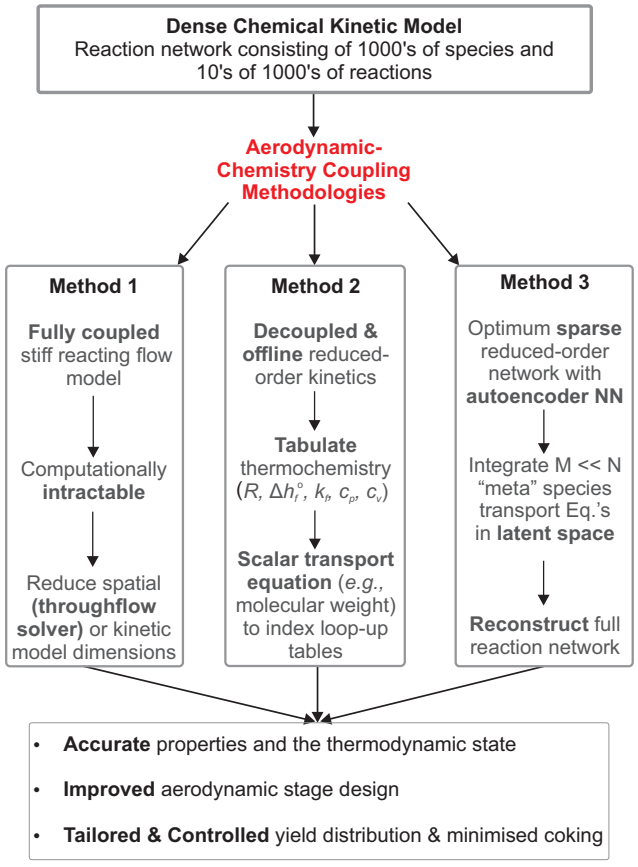


FIGURE 11: **Block diagram illustrating three aerodynamic-chemistry coupling methodologies**

Copyright © 2022 by ASME

Method 1 This high-fidelity methodology enables accurate yield estimates using a fully reacting flow model coupled into a baseline aerodynamic solver. The challenge is that the detailed reaction mechanism for thermal hydrocarbon cracking of high-carbon-number feeds (e.g., naphtha) contains tens of thousands of reactions between thousands of species [44]. Even a simplified reaction network (kinetic model) consists of 20 – 25 additional species transport equations and 20 – 30 stiff chemical reaction equations [45]. Therefore, without intelligent simplification of either the level of spatial dimensionality (*i.e,* using a 2D through-flow solver) or the complexity of the system of reaction equations (see Method 3), the problem is intractable.

Method 2 The second method—and the approach illustrated as an example in this paper—is a low-order framework that entirely decouples the chemical kinetics and aerothermal models. In this approach, the kinetic model is reduced to a single primary (dominant) reaction [45] computed offline, as shown in Eq. 1. The molar concentrations of the chemical species and the thermochemistry (reaction rates, enthalpy of formation, gas constant, specific heat capacity, etc.) of the gas mixture are tabulated within a look-up table as a function of a scalar reaction progress variable (or equivalently molecular weight) and temperature using Aspen [46]. These variables are then transported in the full 3D CFD and used to access/index the tabulated data during each CFD time-step. The scalar transport equation is formed from standard convection and diffusion terms, in addition to a new source term capturing the rate of reaction progress. This methodology is similar to the flamelet generated manifold technique applied in combustion modeling [47].

$$\begin{aligned} C_{6.5}H_{14} \longrightarrow\ & 0.5\,H_2 + 0.76\,CH_4 + 1.16\,C_2H_4 + 0.13\,C_2H_6 \\ & + 0.38\,C_3H_6 + 0.09\,C_3H_8 + 0.008\,C_4H_{10} \\ & + 0.245\,C_4H_8 + 0.113\,C_4H_6 + 0.08\,C_4 \end{aligned} \tag{1}$$

The progress through the primary reaction defined in Eq. 1 is critically dependent on the residence time distribution and flow history (especially temperature) for an average fluid parcel. In order to capture these effects, the progress of reaction ϕ is formulated as a standard scalar transport equation solved within the 3D CFD solver. This equation can be expressed as follows:

$$\frac{\mathrm{D}\rho\phi}{\mathrm{D}t} = \frac{\partial}{\partial x_j}\left(\Gamma\frac{\partial\phi}{\partial x_j}\right) + \underbrace{S_\phi(T)} \tag{2}$$

where $\frac{\mathrm{D}(\cdot)}{\mathrm{D}t}$ is the material derivative and $S_\phi(T)$ is a source term modeling the *rate* of reaction progress from reactants (a hydrocarbon-steam mixture) to products (primarily ethylene). Based on the offline tabulated data, this source term is fitted to an exponential function as function of static temperature (consistent with the functional form of the Arrhenius rate constant). On each time step t_n the cell static temperature is used to update the cell-value for the source term from the regression fit.

The mass-averaged variation in the reaction progress is illustrated in Fig. 12(a). This variable provides the designer valuable insight into the reaction evolution (see Fig. 12(a)) with almost no computational. The next step for low-order chemistry modeling is to capture the effect of the endothermic reactions on the enthalpy profile. For this purpose, a sink term is included in the energy equation as highlighted in red in Eq. 3 below:

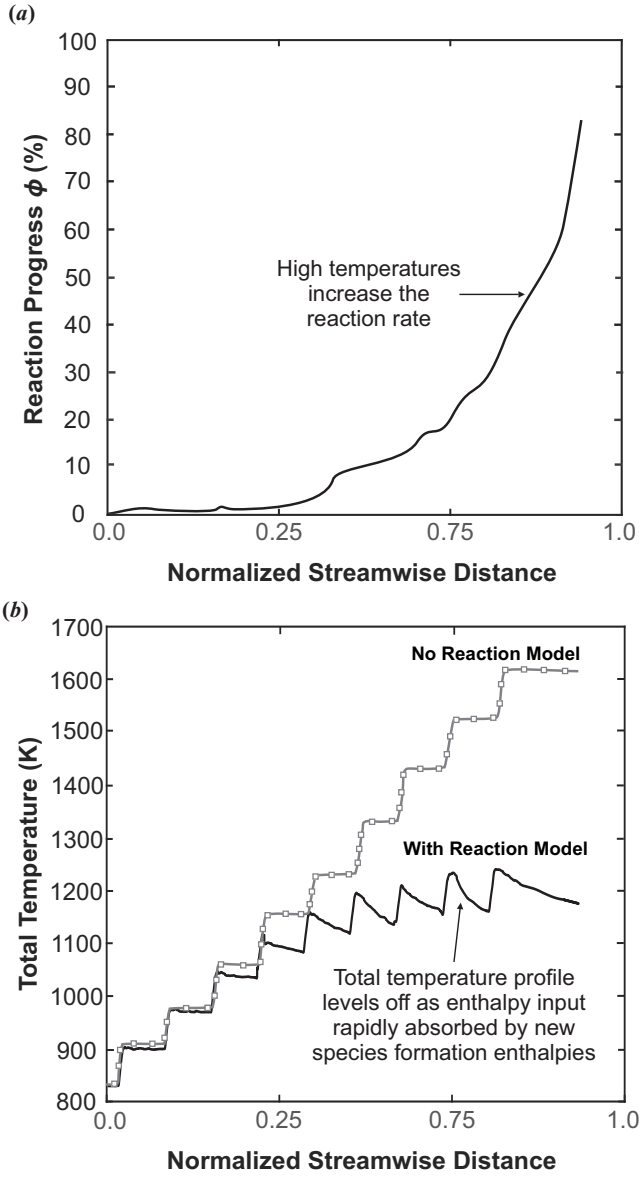


FIGURE 12: **Mass-averaged (a) reaction progress ϕ and (b) total temperature profile along the gas path, both with and without the sink term included in the energy equation**

Copyright © 2022 by ASME

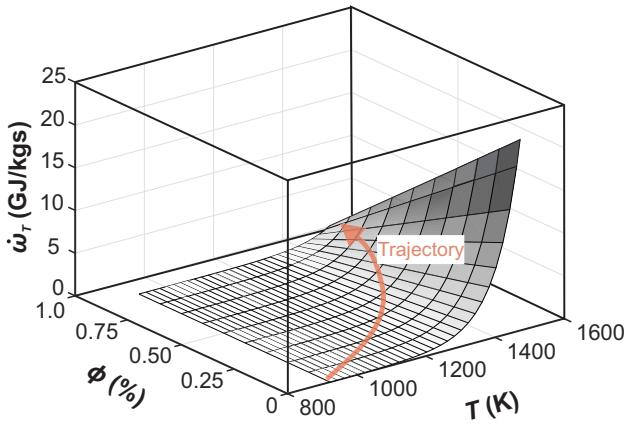


FIGURE 13: **The offline tabulated data for the rate of enthalpy absorbed by the formation enthalpies $\dot{\omega}_T(\phi,T)$ of the species as a function of static temperature T and reaction progress ϕ**

$$\frac{\mathrm{D}\rho E}{\mathrm{D}t} = \frac{\partial}{\partial x_j}\left(\lambda\frac{\partial T}{\partial x_j}\right) + \frac{\partial}{\partial x_j}(\sigma_{ij}u_i) + \dot{Q} + \underbrace{\dot{\omega}_T(\phi,T)}_{} \tag{3}$$

where E is the total *non-chemical* energy (see Eq. 4) and $\dot{\omega}_T(\phi,T)$ is the rate of formation enthalpy absorption. Defining total energy only in terms of sensible energy and kinetic energy—in other words excluding the chemical enthalpy component (which is defined as $E_{\text{chem}} = \sum_{m=1}^{M} \cdot \Delta h^0_{f,m} \cdot Y_m$)—simplifies the required modifications to the energy equation. For this reacting flow case, the only change to non-reacting-flow energy equation is an additional source term $\dot{\omega}_T(\phi,T)$ that captures the lumped effects of the endothermic reactions. This source term $\dot{\omega}_T(\phi,T)$ is defined follows in Eq. 5.

$$E = \int_{T_{\text{ref}}}^{T} C_v(\phi,T)dT - RT_{\text{ref}}/W + \frac{1}{2}u_i u_i \tag{4}$$

$$\dot{\omega}_T(\phi,T) = \sum_{m=1}^{M} \Delta h^0_{f,m}(T)\cdot\dot{\omega}_m(T) \tag{5}$$

In Eq. 5 the summation is over the M species contained within the gas mixture and the summand is the product of the enthalpy of reaction $\Delta h^0_{f,m}(T)$ and the associated volumetric molar production/destruction rate $\dot{\omega}_m$ of species m. The source term $\dot{\omega}_T(\phi,T)$ is tabulated as a function of static temperature T and reaction progress ϕ (see Fig. 13). It is noted that the reaction rate pressure dependence is assumed negligible at this time. This is a valid assumption for low order aerodynamic-chemistry modeling [48].

With the sink term activated, Fig. 12(b) illustrates that the total temperature profile levels-off in the latter stages of the turbo-reactor. In this high-reaction-rate region, the imparted enthalpy is rapidly absorbed in the form of the formation enthalpies of the new species (rather than sensible enthalpy).

The key challenge with this approach is that (1) temperature and reaction progress profiles are not yet accurate and (2) and the yield and coke deposition rates are currently unavailable with Method 2. This is because an insufficient number of species and secondary reactions are accounted for, and transport phenomena are not incorporated. To address this, more reactions could be added. However, it is unclear which reactions to include and more importantly, a clear and distinct single reaction progress variable would no longer available if more reactions were added. Method 3 (below) could potentially address these challenges by identifying the optimum dimensions of the look-up defining the thermochemistry data by learning a reduced-order reaction network.

Method 3 This final method aims to reduce the dimensions (and hence computational cost) of (1) the fluid dynamics (using a throughflow solver), (2) the reaction network or (3) a combination of above. In the first case, Fig. 14 illustrates, as an example, how a dimensionality reduction technique—using a physics-driven autoencoder neural network [49]—can be used to learn the optimal sparse reaction network within a latent space. The number of physical species transport equations can be compressed to only a few "meta" species transport equation in latent space. In order have a generalizable reduced-dimensionality kinetic model, the latent space is constrained by the physics of a true Arrhenius-based "meta" chemical reaction system [49].

In order to generate training data, the design space can be widely and uniformly sampled by perturbing the temperature, pressure, enthalpy input rate, etc. profiles within a 1-D plug flow reactor (PFR) model. The full kinetic model (*i.e.,* up to tens of thousands of reactions) is applied in the PFR, enabling the response of the full reaction network to be explored. Since this data can be generated at low cost using this 1-D PFR approach, Fig. 14 (left) demonstrates that a large number of training data samples can be generated for the neural network. In order to automatically generate the reaction network RMG-PY is used [50]. Following network construction, the ODE stiff system of equations is then simulated using a PFR within Cantera [51].

A key feature of this autoencoder architecture is the decoder. This enables the full species concentration distributions to always be reconstructed/recovered from latent space, enabling accurate yield and coking rate predictions once steady state is reached.

This low-order reaction network can then either by coupled into a 3-D CFD or 2-D throughflow solver. The throughflow solver is an ideal candidate for low-order reacting flow modeling especially when performing multiobjective design optimization where thousands of different design conditions must be simulated (as discussed at the start of Sec. 4.1 & 5.1) in order to tailor the reaction product distribution for different requirements whilst minimizing coke deposition rate. This is a particularly attractive

Copyright © 2022 by ASME

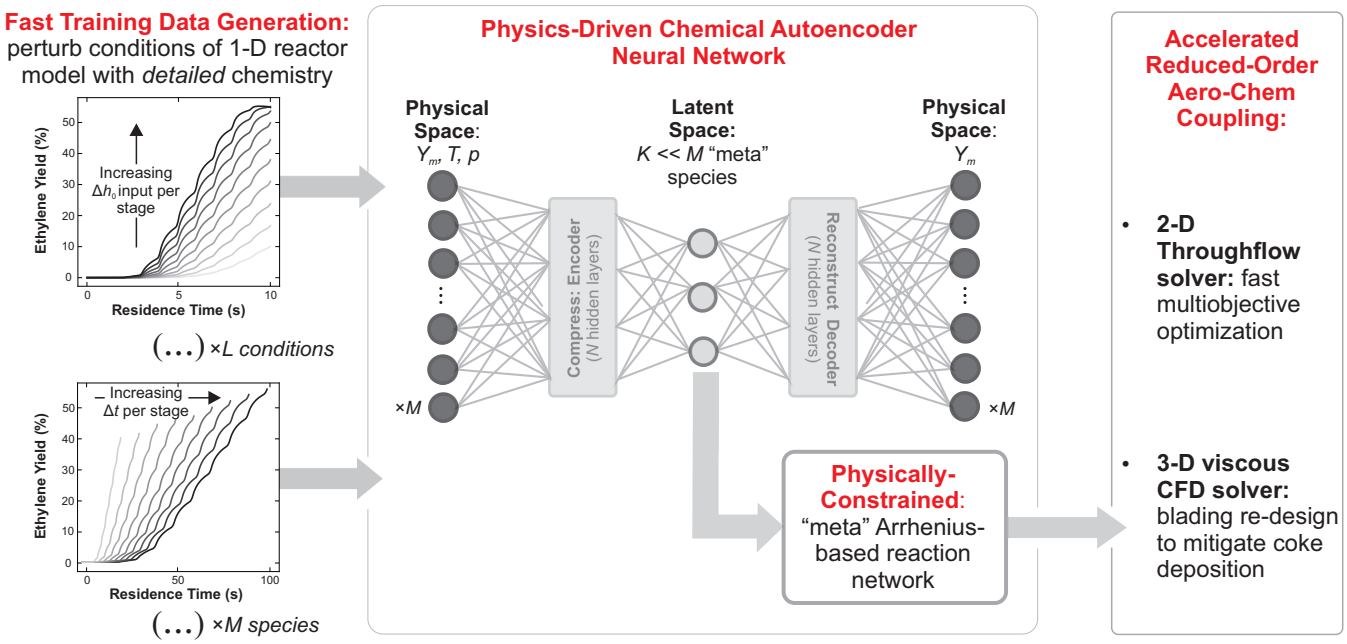


FIGURE 14: **Chemical autoencoder neural network for reduced-dimensionality aerodynamic-chemistry coupling methodology**

solution if the throughflow solver has been corrected with LES calibration data, as described in Sec. 4.1. Further details of this methodology (Method 3) will be presented in future publications.

6 SUMMARY AND CONCLUSIONS

(1) This paper has introduced a new concept based on an ultra-high power-density turbomachine. The objective of the turbo-reactor concept is to decarbonize the carbon-intensive hydrocarbon cracking process via direct renewable electrification. The urgent demand to decarbonize necessitates an accelerated design-cycle time frame for the new concept.

(2) It has been highlighted that the novel design requirements and working fluid properties have enabled an ultra-high power density machine architecture. Consequently, the level of complexity of the flow physics is elevated. These associated fluid dynamic interactions include: low Reynolds number large-scale flow separation and blockage, highly supersonic flow, strong three-dimensionality, shockwave boundary layer interaction, intense turbulent mixing, and a low Reynolds number transitional flow regime with extended regions of non-equilibrium turbulence.

(3) As a result, there is limited manufacturer experience and experimental data available to inform and validate the early design process as well as calibrate low-order numerical tools. The complexity of the flow physics, coupled with a lack of suitable empirical correlations, has led to uncertainty in the predictive capabilities of low-order numerical tools.

(4) This paper emphasizes the importance of a tailored design ecosystem aimed to address the aforementioned limitations within the design cycle. Whilst the relatively low Reynolds number flow regime may be a key driver promoting the complexity of the flow physics, it is also the principal enabler for (i) routine and upfront LES during both preliminary and detailed design phases and (ii) providing the robustness of the design. It is demonstrated that this high-fidelity CFD data can be used to calibrate blockage, deviation and loss correlations within throughflow solvers as well as tune turbulence closure models within RANS solvers. Since the off-design amplification of incidence and deviation across the multi-stage machine is eliminated, it has been emphasized that the design is desensitized to modeling uncertainties.

(5) A key feature of the new design system is an aerodynamic-chemistry coupling framework. It has been demonstrated that an offline, tabulated kinetic model approach enables vital information to be incorporated into 3D aerodynamic simulations in an efficient manner. This ensures that (i) the CFD simulations account for the large variation in fluid properties due to the reaction, and (ii) early within the design phase, the aerodynamic designer is provided with insight into the progress of the reaction and the associated impact on the enthalpy distribution. Therefore, the designer can quickly tailor geometric features, *e.g.,* the vaneless space length, to boost the reaction product yield. In order to per-

Copyright © 2022 by ASME

form yield and coking rate calculations for design optimization in a 2-D or 3-D computational environment, a physics-driven chemical autoencoder neural network methodology has been introduced. It has been demonstrated that broad range of training data samples spanning the design space can be generated rapidly by perturbing a 1-D PFR incorporating a detailed kinetic model.

ACKNOWLEDGMENT

The authors would like to thank Coolbrook Oy for supporting this work and permitting the publication of this study. The authors also thank Dr. Tuomas Ouni for providing the thermochemistry data used within this work. The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing facility in carrying out this work.

References

[1] European Commission: Secretariat General, 2021. "Fit for 55: delivering the EU's 2030 Climate Target on the way to climate neutrality".

[2] Seppl, J., Hiltunen, J., and Purola, V.-M., 2017. Process and rotary machine type reactor. *European Patent Specification No. EP 3 164 207 B1*. Munich, European Patent Office.

[3] Coolbrook Oy, 2020. Coolbrook Oy. Helsinki, Finland, accessed Oct. 15, 2020, https://coolbrook.com.

[4] Gross, S., 2021. "The challenge of decarbonizing heavy industry". *Brookings India*.

[5] Roelofsen, O., Somers, K., Speelman, E., and Witteveen, M., 2020. "Plugging in: What electrification can do for industry". *McKinsey & Company. May,* ***28***.

[6] Rubini, D., Xu, L., Rosic, B., and Johannesdahl, H., 2021. "A New Turbomachine for Clean and Sustainable Hydrocarbon Cracking". *ASME J. Eng Gas Turbines Power,* ***144***(2), 12. 021024.

[7] Rubini, D., Karefyllidis, N., Xu, L., Rosic, B., and Johannesdahl, H. "A New Robust Regenerative Turbo-Reactor Concept for Clean Hydrocarbon Cracking". *Journal of the Global Power and Propulsion Society,* ***6***, pp. 1–16.

[8] Mayle, R. E., 1991. "The Role of Laminar-Turbulent Transition in Gas Turbine Engines". *Volume 5: Manufacturing Materials and Metallurgy; Ceramics; Structures and Dynamics; Controls, Diagnostics and Instrumentation; Education; IGTI Scholar Award; General. Orlando, Florida, USA. June 3–6*. V005T17A001.

[9] Liu, Q., Ager, W., Hall, C., and Wheeler, A. P., 2022. "Low Reynolds Number Effects on the Separation and Wake of a Compressor Blade". *ASME J. Turbomach,* ***144***(10), 04. 101008.

[10] Craig, H., and Cox, H., 1970. "Performance estimation of axial flow turbines". *Proceedings of the Institution of Mechanical Engineers,* ***185***(1), pp. 407–424.

[11] Lieblein, S., 1965. Aerodynamic Design of Axial-flow Compressors. VI - Experimental Flow in Two-Dimensional Cascades. Research memorandum, National Advisory Committee for Aeronautics, Washington.

[12] Stratford, B., and Sansome, G., 1960. "Theory and tunnel tests of rotor blades for supersonic turbines". *Ministry of Aviation: Aeronautical Research Council Reports and Memoranda*.

[13] Bufi, E., Cinnella, P., Obert, B., et al., 2015. "Fast design methodology for supersonic turbines with strong real gas effects". In Congrès français de mécanique, AFM, Association Française de Mécanique.

[14] Cinnella, P., and Elio, A. B., 2017. "Design methodology for supersonic ORC turbine blades". In CFM 2017-23ème Congrès Français de Mécanique, AFM, Maison de la Mécanique, 39/41 rue Louis Blanc-92400 Courbevoie.

[15] Sixsmith, H., and Altmann, H., 1977. "A Regenerative Compressor". *Journal of Engineering for Industry,* ***99***(3), 08, pp. 637–647.

[16] Liccini, L. L., 1949. "Analytical and experimental investigation of 90° supersonic turning passages suitable for supersonic compressors or turbines". *NACA RM L9G07*.

[17] Colclough, C., 1966. "Design of Turbine Blades Suitable for Supersonic Relative Inlet Velocities and the Investigation of Their Performance in Cascades: Part II-Experiments, Results and Discussion". *Journal of mechanical engineering science,* ***8***(2), pp. 185–197.

[18] Power, B., Xu, L., and Wellborn, S., 2014. "Numerical and Experimental Findings of a Highly-Loaded Aspirated Cascade". *Proceedings of the ASME Turbo Expo 2014: Turbine Technical Conference and Exposition. Volume 2A: Turbomachinery. Düsseldorf, Germany. June 16–20*. V02AT37A050.

[19] Dickens, T., and Day, I., 2010. "The Design of Highly Loaded Axial Compressors". *ASME J. Turbomach,* ***133***(3), 11. 031007.

[20] Chana, K. S., and Miller, R. J., 2020. "The Effect of Reaction on Compressor Performance". *Proceedings of the ASME Turbo Expo 2020: Turbomachinery Technical Conference and Exposition. Volume 2A: Turbomachinery. Virtual, Online. September 21–25*. V02AT32A047.

[21] Spurr, A., 1980. "The prediction of 3D transonic flow in turbomachinery using a combined throughflow and blade-to-blade time marching method". *International Journal of Heat and Fluid Flow,* ***2***(4), pp. 189–199.

[22] Denton, J. D., 1978. "Throughflow Calculations for Transonic Axial Flow Turbines". *ASME J. Eng. Power,* ***100***(2), 04, pp. 212–218.

[23] Greitzer, E. M., 1981. "The Stability of Pumping Systems—The 1980 Freeman Scholar Lecture". *ASME J. Fluids Eng,* ***103***(2), 06, pp. 193–242.

[24] Pullan, G., Young, A. M., Day, I. J., Greitzer, E. M., and Spakovszky, Z. S., 2015. "Origins and Structure of Spike-

Copyright © 2022 by ASME

Type Rotating Stall". *ASME J. Turbomach,* ***137***(5), 05. 051007.

[25] Arroyo, C. P., Dombard, J., Duchaine, F., Gicquel, L., Martin, B., Odier, N., and Staffelbach, G., 2021. "Towards the Large-Eddy Simulation of a full engine: Integration of a 360 azimuthal degrees fan, compressor and combustion chamber. Part I: Methodology and initialisation". *Journal of the Global Power and Propulsion Society,* ***2021***(May), pp. 1–16.

[26] Tyacke, J., Tucker, P., Jefferson-Loveday, R., Vadlamani, N. R., Watson, R., Naqavi, I., and Yang, X., 2013. "LES for Turbines: Methodologies, Cost and Future Outlooks". *Proceedings of the ASME Turbo Expo 2013: Turbine Technical Conference and Exposition. Volume 6B: Turbomachinery. San Antonio, Texas, USA. June 3–7.* V06BT37A007.

[27] Shaikh, F., and Rosic, B., 2021. "Unsteady phenomena at the combustor-turbine interface". *Journal of the Global Power and Propulsion Society,* ***5***, pp. 202–215.

[28] Jacobi, S., Mazzoni, C., Rosic, B., and Chana, K., 2017. "Investigation of unsteady flow phenomena in first vane caused by combustor flow with swirl". *ASME J. Turbomach,* ***139***(4), p. 041006.

[29] Miki, K., and Ameri, A., 2021. "Improved Prediction of Losses with Large Eddy Simulation in a Low- Pressure Turbine". *ASME J. Turbomachinery*, 12, pp. 1–38.

[30] Tyacke, J. C., and Tucker, P. G., 2015. "Future use of large eddy simulation in aero-engines". *ASME J. Turbomach,* ***137***(8), p. 081005.

[31] Choi, H., and Moin, P., 2012. "Grid-point requirements for large eddy simulation: Chapman's estimates revisited". *Physics of fluids,* ***24***(1), p. 011702.

[32] Denton, J., and Dawes, W., 1998. "Computational fluid dynamics for turbomachinery design". *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science,* ***213***(2), pp. 107–124.

[33] Banjac, M., and Petrovic, M. V., 2018. "Development of method and computer program for multistage axial compressor design: Part I—Mean line design and example cases". In *Turbo Expo: Power for Land, Sea, and Air*, Vol. 51012, American Society of Mechanical Engineers, p. V02CT42A007.

[34] Banjac, M., Petrovic, M. V., and Wiedermann, A., 2014. "A new loss and deviation model for axial compressor inlet guide vanes". *ASME J. Turbomachinery,* ***136***(7), p. 071011.

[35] Pasquale, D., Persico, G., and Rebay, S., 2013. "Optimization of Turbomachinery Flow Surfaces Applying a CFD-Based Throughflow Method". *ASME J. Turbomach.,* ***136***(3), 09. 031013.

[36] Spalart, P., and Allmaras, S., 1992. "A one-equation turbulence model for aerodynamic flows". In 30th aerospace sciences meeting and exhibit, p. 439.

[37] Spalart, P. R., 1997. "Comments on the feasibility of les for wings, and on a hybrid rans/les approach". In Proceedings of first AFOSR international conference on DNS/LES, Greyden Press.

[38] Menter, F., and Egorov, Y., 2010. "The scale-adaptive simulation method for unsteady turbulent flow predictions. part 1: theory and model description". *Flow, turbulence and combustion,* ***85***(1), pp. 113–138.

[39] He, X., Zhao, F., and Vahdati, M., 2021. "Uncertainty Quantification of Spalart–Allmaras Turbulence Model Coefficients for Compressor Stall". *ASME J. Turbomachinery,* ***143***(8), 04. 081007.

[40] Schnoes, M., and Nicke, E., 2017. "Exploring a database of optimal airfoils for axial compressor design".

[41] Mudalige, G., Giles, M., Reguly, I., Bertolli, C., and Kelly, P., 2012. "OP2: An active library framework for solving unstructured mesh-based applications on multi-core and many-core architectures". In 2012 Innovative Parallel Computing (InPar), pp. 1–12.

[42] Denton, J. D., 2017. "Multall—an open source, computational fluid dynamics based, turbomachinery design system". *ASME J. Turbomach,* ***139***(12), p. 121001.

[43] Denton, J. D., 1992. "The calculation of three-dimensional viscous flow through multistage turbomachines". *ASME J. Turbomach,* ***114***(1), 01, pp. 18–26.

[44] Heynderickx, G. J., and Froment, G. F., 1996. "A pyrolysis furnace with reactor tubes of elliptical cross section". *Industrial & engineering chemistry research,* ***35***(7), pp. 2183–2189.

[45] Kumar, P., and Kunzru, D., 1985. "Kinetics of coke deposition in naphtha pyrolysis". *The Canadian Journal of Chemical Engineering,* ***63***(4), pp. 598–604.

[46] Al-Malah, K. I., 2016. *Aspen plus: chemical engineering applications*. John Wiley & Sons.

[47] Klapdor, E. V., 2011. "Simulation of combustor-turbine interaction in a jet engine". PhD thesis, Technische Universität Darmstadt.

[48] Van Geem, K. M., Reyniers, M.-F., Marin, G. B., Song, J., Green, W. H., and Matheu, D. M., 2006. "Automatic reaction network generation using rmg for steam cracking of n-hexane". *AIChE Journal,* ***52***(2), pp. 718–730.

[49] Grassi, T., Nauman, F., Ramsey, J., Bovino, S., Picogna, G., and Ercolano, B., 2021. "Reducing the complexity of chemical networks via interpretable autoencoders". *arXiv preprint arXiv:2104.09516*.

[50] Gao, C. W., Allen, J. W., Green, W. H., and West, R. H., 2016. "Reaction mechanism generator: Automatic construction of chemical kinetic mechanisms". *Computer Physics Communications,* ***203***, pp. 212–225.

[51] Goodwin, D. G., Moffat, H. K., Schoegl, I., Speth, R. L., and Weber, B. W., 2022. Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes. https://www.cantera.org. Version 2.6.0.

Copyright © 2022 by ASME

Statement of Authorship for joint/multi-authored papers for PGR thesis

To appear at the end of each thesis chapter submitted as an article/paper

The statement shall describe the candidate's and co-authors' independent research contributions in the thesis publications. For each publication there should exist a complete statement that is to be filled out and signed by the candidate and supervisor **(only required where there isn't already a statement of contribution within the paper itself).**

Title of Paper	Accelerating the Development of a New Turbomachinery Concept in an Environment With Limited Resources and Experimental Data: Challenges
Publication Status	□Published □ Accepted for Publication □Submitted for Publication □Unpublished and unsubmitted work written in a manuscript style
Publication Details	Rubini, D, Karefyllidis, N, Xu, L, Rosic, B, & Johannesdahl, H. "Accelerating the Development of a New Turbomachinery Concept in an Environment With Limited Resources and Experimental Data: Challenges." *Proceedings of the ASME Turbo Expo 2022: Turbomachinery Technical Conference and Exposition*. *Volume 10D: Turbomachinery — Multidisciplinary Design Approaches, Optimization, and Uncertainty Quantification; Turbomachinery General Interest; Unsteady Flows in Turbomachinery*. Rotterdam, Netherlands. June 13–17, 2022. V10DT36A001. ASME. https://doi.org/10.1115/GT2022-80698

Student Confirmation

Student Name:	Dylan Rubini		
Contribution to the Paper	Performed majority of numerical simulations, developed required codes for simulation and postprocessing, post-processed and analyzed results and wrote manuscript		
Signature		Date	16/08/2024

Supervisor Confirmation

By signing the Statement of Authorship, you are certifying that the candidate made a substantial contribution to the publication, and that the description described above is accurate.

Supervisor name and title: Prof. Budimir Rosic			
Supervisor comments I confirm that the candidate's statement is an accurate description of his contribution to this publication.			
Signature		Date	19/08/2024

This completed form should be included in the thesis, at the end of the relevant chapter.

5.1.1 Summary

In this chapter, the unique level of complexity of flow physics within the turbo-reactor has been illustrated. As a result, existing turbomachinery experimental databases are not sufficiently accurate, and standard RANS turbulence models exhibit elevated uncertainties. High-fidelity numerical simulations have been used to remedy this problem. However, this does not solve the computational challenges of coupling chemical kinetics to CFD simulations. A highly simplified method of aerochemical coupling has been proposed. This uses only a single reaction equation in the kinetic model and therefore cannot provide an estimate of the product yields or coking rates, which are the ultimate design metrics. To improve the bladed-path aerodynamic design for maximum yield and minimum coking rate, this simplified method needs to be improved. Therefore, we are motivated to develop a new accelerated method for aerothermochemical coupling. This method is introduced and verified in the next three chapters. These form "Part 2" of this thesis, which brings together aerothermodynamics and chemistry.

Part II

Aerothermochemistry

Chapter 6

Modelling Aerochemical Interactions

6.1 Introduction

Aerothermochemical optimisation of the bladed flow path presents an entirely new requirement for turbomachinery design. As eluded to in Chapter 3, the aerodynamic design must be optimised to improve the chemical reaction performance. Since feedstock costs outweigh energy costs (see Fig. 1.5), relatively small improvements in reaction performance could make electrification cost competitive with gas furnaces. However, although it is vital for the success of the turbo-reactor concept, aerothermochemical optimisation has so far been left virtually unexplored because of the prohibitive computational burden of numerical reacting flow simulations using realistic, detailed and comprehensive chemical kinetic models (such as Refs. [88, 94–97]). Even a reacting flow simulation with a small kinetic model and a coarse mesh can take on the order of days to run. Now imagine querying the objective function hundreds to thousands of times in an outer optimisation loop. This would take years, which by definition is the meaning of intractability! For heavy feedstocks such as naphtha, the computational cost is further amplified as the number of potential reaction pathways grows non-linearly with carbon number.

Therefore, the numerical design and analysis of the turbo-reactor for realistic feeds and mechanisms have so far been restricted to a decoupled aerochemical modelling approach, where information about the yield and coking rates (see Appendix A) are determined solely from simple, idealised, homogeneous 1-D chemical kinetic calculations (see Fig. 4.1 in Sec. 4.1.1). While this gives realistic estimates for simple flows like tubular chemical furnaces [12], the same cannot be said for the turbo-reactor. Therefore, more accurate predictions of basic performance metrics are available only after costly and time-consuming physical experiments—and these are still limited by low spatial resolution and so cannot guide design optimisation.

Despite the increase in computing power with Moore's law [98], it is still not feasible to incorporate more than 10 to 20 species at the industrial design level [12]. This simplification can introduce considerable uncertainties—particularly for minor species and regions of over-cracking—as shown using chemical kinetic simulations in Fig. 6.1. Although the simplified chemistry is 70 times faster (in this example) than the detailed model, the errors are not tolerable. The errors are primary because simple mechanisms are often only valid over a limited operating range, prohibiting design space exploration. Therefore, it is desirable to use detailed and comprehensive mechanisms [99].

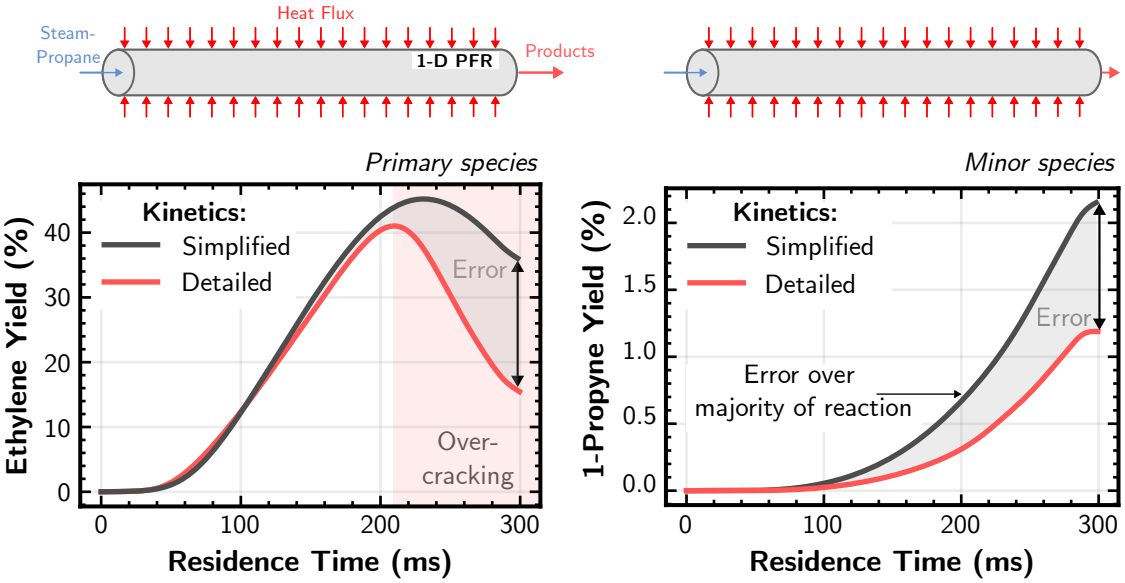


Figure 6.1: Chemical kinetic calculations of a 1D PFR demonstrating the risks of using a simplified (small) kinetic model compared to a detailed, comprehensive and accurate (large) model

The aforementioned limitations of the current design system can lead to disastrously sub-optimal reaction efficiency measured for the manufactured machine (relative to the design intent) and result in expensive iterative re-design and re-manufacturing. Given the accelerating climate emergency and the urgency for the net-zero transition, this laborious process is not sustainable. To address this, the following chapters will introduce an efficiently coupled multi-fidelity machine-learning-assisted (ML) aerochemical modelling methodology, called `ChemZIP`, which dramatically accelerates complex multi-species reacting flow calculations by two to three orders of magnitude, while maintaining training times within a few hours. This methodology captures the dynamic response of a

detailed and realistic kinetic model (>200 species), while retaining a high-fidelity industry-standard 3-D viscous fluid model and without imposing additional time step restrictions on the basic solver. The `ChemZIP` methodology will enable aerothermochemically-guided computational optimisation to be applied routinely, for the first time, in the early phases of the design cycle. This will enable aerodynamic design for a more efficient reaction, reducing energy consumption and increasing the amount of high-value product produced for a given mass of feed.

Before introducing this methodology, the importance of capturing aerothermochemical interactions in the design process of the turbo-reactor will be explained and detailed. This will be accompanied by a brief review of state-of-the-art approaches to accelerating reacting flow simulations, highlighting their limitations. Finally, the new ML-assisted `ChemZIP` methodology will be introduced in Chapter 7, followed by a detailed verification comparison against a state-of-the-art commercial solver using detailed chemical kinetics in Chapter 8. To prove the viability of the new methodology in a turbomachinery setting, `ChemZIP` is integrated within an industry-standard 3D viscous turbomachinery fluid solver (*i.e.,* TBLOCK). The results provide the first numerical predictions of the yield and coking rate distributions along the bladed path of the turbo-reactor. Although similar data-driven acceleration approaches have been exploited in the last few years in the field of combustion modelling, they have never been applied to aerochemical processes along the bladed path of turbomachines, and more generally, there is surprisingly limited focus on the chemical process industry.

6.1.1 Importance of 3-D Viscous Aerochemical Interactions

It is clear that there are three primary design objectives for the turbo-reactor: maximising reaction efficiency and selectivity by designing a reaction-efficient temperature profile (see Chapter 3), minimising coke deposition on the metal surfaces, and maximising power density whilst ensuring stability, durability and availability. Although the latter is possible to achieve effectively with purely aerothermal computational design and modelling, the former two require affordable 2-D/3-D viscous reacting flow simulations

to be realised in practise. This is because of important interactions between the aerodynamics and the chemistry.

When the spatio-temporal scales of chemistry and aerodynamics are similar, the two phenomena can interact bidirectionally. Figure 6.2 indicates scale similarity over a wide range of fluid time-scales, from the global residence time-scale down to the Kolmogorov time-scale. This means that the chemistry is in the *non-equilibrium finite-rate* kinetics regime and aerothermochemical interactions are important. The degree of aerochemical coupling can be succinctly characterised by the Damköhler [100]

$$Da_{\mathrm{I},k} = \frac{\tau_f}{\tau_{c_k}} \tag{6.1}$$

where τ_f is the characteristic fluid time-scale and τ_{c_k} is the characteristic chemical time-scale of species k (see Eqs. (6.12) & (6.13) in Sec. 6.2). In the limit of $Da_{\mathrm{I},k} \to \infty$ the reaction equilibrates infinitely fast, and in the limit of $Da_{\mathrm{I},k} \to 0$ the chemistry is frozen. However, in the turbo-reactor, $Da_{\mathrm{I},k}$ is closer to 1 (see Fig. 6.2) so interactions are important. The following sections identify some key features of these interactions and discuss the significance of one-way and two-way aerochemical coupling.

Fluid → Chemistry Interactions

The aerodynamic environment in the turbo-reactor is highly complex (see Chapter 5). This results in non-uniformities, turbulent mixing and/or molecular diffusion on a scale to which the reaction responds appreciably (see Fig. 6.2). Complex two-/three-dimensional viscous processes, including shocks, are not accounted for by 1-D decoupled idealised kinetic calculations (see Sec. 4.1.1), and therefore numerical predictions of the yield and coking rate distributions are unreliable. Furthermore, complex three-dimensional flow physics such as interpass mixing—a unique feature of the regenerative toroidal turbo-reactor architecture (see Section 2.2)—cannot be quantified without 3-D viscous reacting flow analysis. These complex multi-dimensional transport phenomena natural to turbomachines amplify the importance of detailed reacting flow simulations compared with tubular furnaces. In reaction coils, the flowfield is simplified and more uniform, allowing the 3D domain to be collapsed into a one-dimensional model without a substantial loss of accuracy.

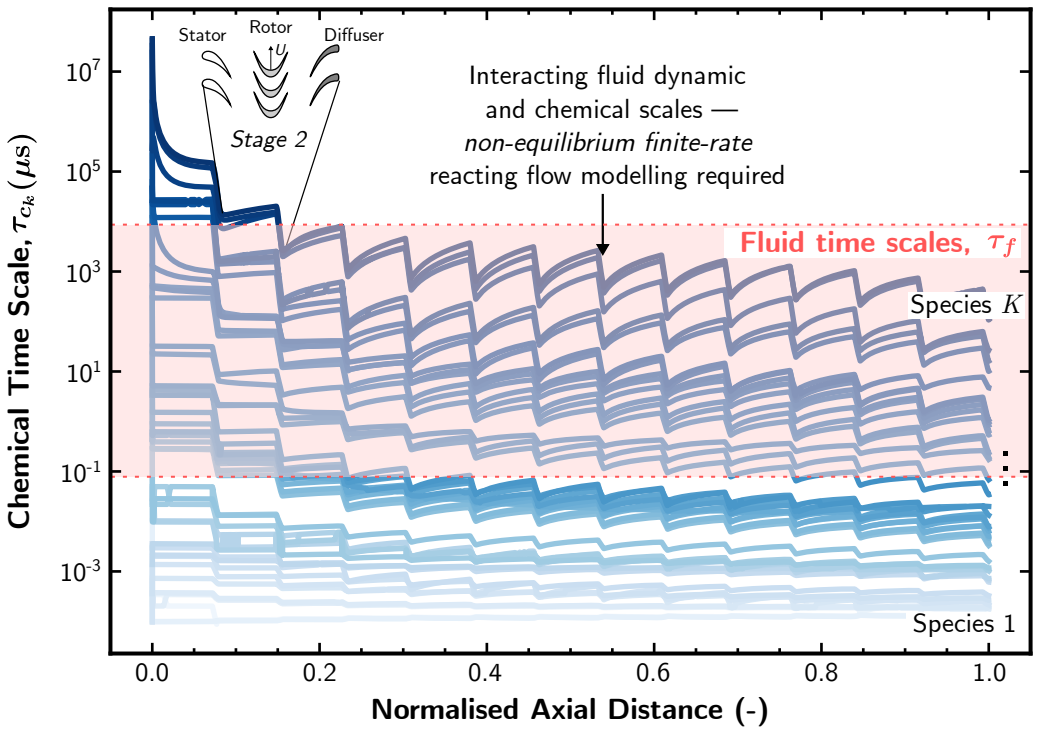


Figure 6.2: Streamwise distribution of characteristic chemical time-scales plotted on a log-scale for an propane feed calculated from 1-D PFR chemical kinetic calculations. The characteristic chemical time scales are calculated using method outlined in Ref. [101]. Time-scales longer than 1000 s (*i.e.,* non-reacting species) are eliminated

While it is apparent that the complex turbomachinery aerothermodynamics can have a notable influence on the chemistry (*i.e.,* fluid → chemistry interactions), the importance of two-way interactions is unclear.

Chemistry → Fluid Interactions

Generally, aerothermochemical interactions are bidirectional, featuring non-linear feedback mechanisms. There are two key chemistry → fluid interactions:

1. Thermophysical fluid properties such as isobaric heat capacity, molecular viscosity, specific gas constant, etc., are composition-dependent; however, this dependency is weak.

2. The endothermic heat of reaction added to the energy equation considerably changes the temperature field, which in turn induces changes in density and fluid properties, accelerating/decelerating the flow.

Heat absorption by the reaction sets up non-linear feedback between the reaction and energy equation, which is compounded by the exponential temperature dependence of the reaction rates. However, for conditions and time-scales representative of the turbo-reactor, the momentum equation is not strongly affected. Inviscid 1-D compressible shockwave calculations shown in Fig. 6.3 suggest that the chemistry → fluid interaction is weak for the momentum equation. The velocity distribution downstream of the shockwave is almost indistinguishable from the non-reacting flow solution. This is because, in comparison with combustion, the density changes are not as significant. Therefore, it is primarily the energy equation that mediates thermochemical feedback. For further details on aerochemical interactions, the reader is referred to Oran & Boris [100].

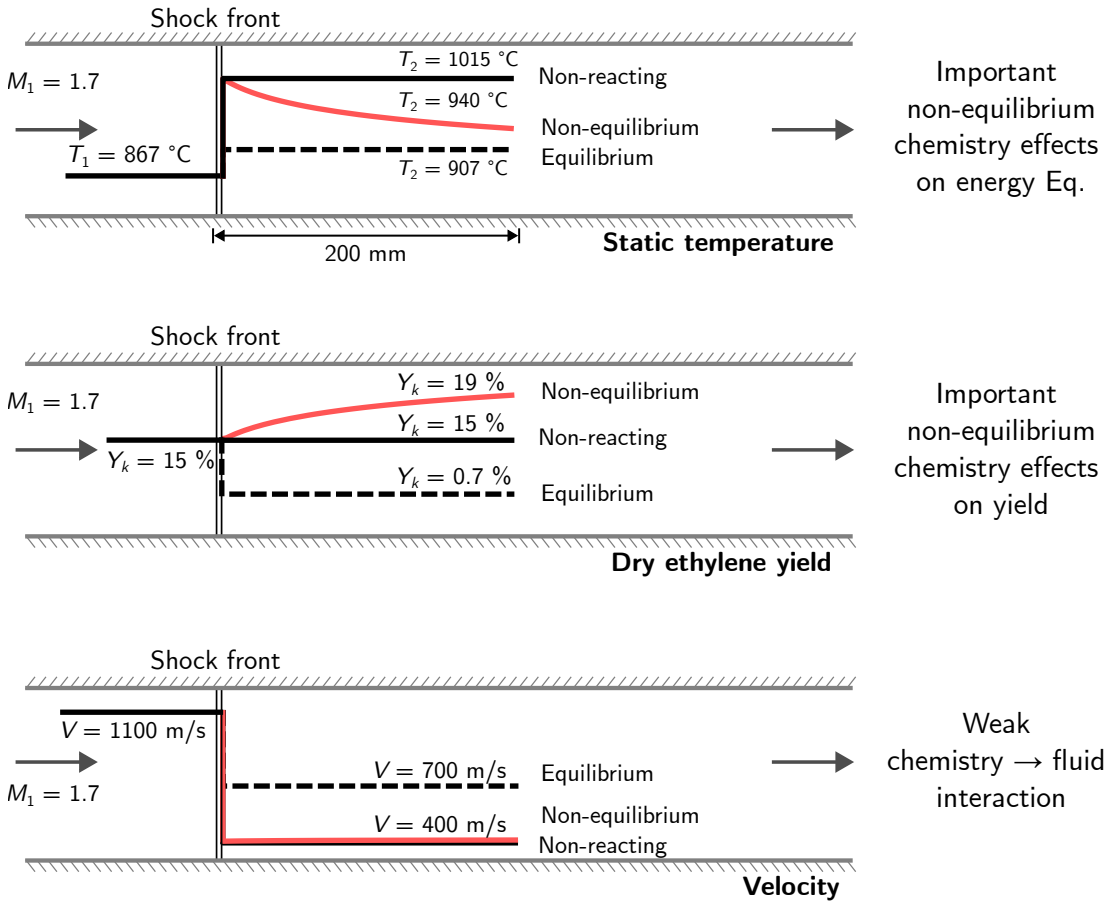


Figure 6.3: Steady, inviscid, adiabatic 1-D compressible flow calculations across a shockwave for a steam-propane gas mixture that is (1) non-reacting (2) in chemical equilibrium and (3) in chemical non-equilibrium. The conditions upstream of the shock are representative of those of the turbo-reactor. The governing equations can be found in Anderson [102]

6.1.2 Challenges of Reacting Flow Modelling

Basic non-reacting CFD alone often requires state-of-the-art high-performance computational resources to solve the five governing equations. Since the chemistry is in non-equilibrium (see Figs. 6.2 & 6.3), a further K_{species} non-linearly coupled species transport equations must be included. This significantly increases computational overheads. For detailed kinetic mechanisms, such as that found in Ref. [88], K_{species} can be greater than 900, or even larger for heavy hydrocarbons. If the chemistry were instead in equilibrium, which is typically the case for hypersonic flows, numerical modelling would be much simpler. Differential equations would not be necessary; instead, a much simpler system of *algebraic* equations would be solved. On top of this, the reactions networks, which can have order 10,000 pathways (see Fig. 6.4), are highly non-linear, tightly coupled and highly multi-scale with time-scales varying by up to 12 orders of magnitude [103] (see Fig. 6.2), down to the sub-nanosecond scale. This intense numerical stiffness further adds to the long convergence time frames.

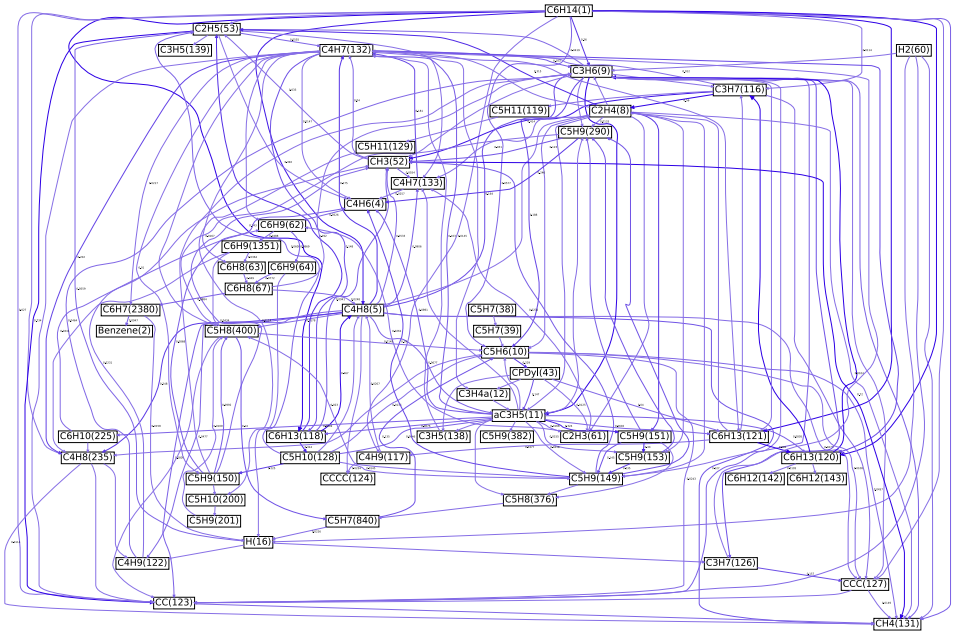


Figure 6.4: Reaction path diagram for a propane pyrolysis chemical mechanism with 65 species and 1700 reactions

The largest computational overhead for non-equilibrium reacting flow modelling comes from having to evaluate the reaction network and construct, LU-factorise and invert the Jacobian matrix, which is necessary for implicit stiff ordinary differential equation (ODE)

solvers [90]. Lu & Law [99] estimated the computational cost of evaluating and factorising the Jacobian scales cubically with the number of species $\mathcal{O}(K^3)$. Consequently, for large mechanisms, the chemical kinetics are by far the bottleneck of a CFD computation, taking over 98% of the CPU time when there are more than 50 species present. For example, direct integration (DI) of a relatively small 65 species mechanism on a 1 million node 3-D mesh takes 20 days on a good 18-core CPU.

For turbomachinery calculations, these problems are exacerbated by the need for larger domain sizes and finer mesh resolutions compared to conventional chemical reactors. Due to the sharper gradients present in the flow, higher grid resolutions are required. Ultimately, all of the challenges of modelling reacting flow mean that it is often prohibitively costly for all but the simplest chemical mechanisms, geometries and/or meshes.

Before reviewing the literature on typical approaches for computational acceleration, we introduce the full-order model for chemically reacting flows. This will help to better understand the challenges and complexity of full-order modelling.

6.2 Fully-Coupled Detailed Reacting Flow Model

This section presents the governing conservation equations for the full-order model (FOM) of multicomponent chemically reacting flow, alongside an explanation and justification of the physical modelling assumptions employed. For in-depth discussions and derivations of the governing equations, the reader is referred to Kee *et al.* [90], Kuo [104], Poinsot & Vaynante [105] and/or Oran & and Boris [100]. The governing equations and physical assumptions presented here form the basis for the `ChemZIP` methodology introduced in Chapter 7.

6.2.1 Conservation Equations

The FOM time-dependent compressible URANS species transport equation is given by

$$\frac{\partial(\rho Y_k)}{\partial t} + \nabla \cdot (\rho \vec{\mathbf{u}} Y_k) = \nabla \cdot \left(\left(\rho \bar{D}_k + \frac{\mu_t}{Sc_t} \right) \nabla Y_k \right) + \rho \dot{\omega}_k \quad \text{for } k = 1, \ldots, K, \tag{6.2}$$

where K is the total number of species, ρ is the density, $\vec{\mathbf{u}}$ is the velocity vector, μ_t is the turbulent viscosity, Sc_t is the turbulent Schmidt number, and for the k^{th} species, $Y_k \in \mathbb{R}^K$ is the species mass fraction, $\bar{D}_k = \bar{D}_k(T, Y_1, \ldots, Y_K)$ is the mixture-averaged Fickian mass diffusion coefficient and $\dot{\omega}_k$ is the net species production rate. The overbar $\bar{\square}$ represents "mixture-averaging", which is a mass-fraction-based weighting of the quantity of interest for each species. The URANS energy equation for the total non-chemical energy

$$E = \underbrace{u_s}_{\text{sensible}} + \underbrace{\frac{1}{2}\vec{\mathbf{u}} \cdot \vec{\mathbf{u}}}_{\text{kinetic}} \tag{6.3}$$

is given by

$$\frac{\partial(\rho E)}{\partial t} + \nabla\cdot(\rho\vec{\mathbf{u}}E) = \nabla\cdot\left(\left(\bar{\lambda} + \frac{c_p\mu_t}{Pr_t}\right)\nabla T\right) + \nabla\cdot(\boldsymbol{\sigma}\cdot\vec{\mathbf{u}}) + \underbrace{\nabla\cdot\left(\rho\sum_{k=1}^{K} h_k\bar{D}_k\nabla Y_k\right)}_{\text{for reacting flow only}} + \dot{\omega}_T, \tag{6.4}$$

where u_s is the sensible component of internal energy, $\bar{\lambda}$ is the molecular thermal conductivity, Pr_t is the turbulent Prandtl number, $\boldsymbol{\sigma}$ is the total stress tensor for a Newtonian fluid, $\dot{\omega}_T$ is the chemical reaction heat absorption/release rate, and h_k is the partial enthalpy of the k^{th} species, defined as

$$h_k = \underbrace{\int_{T^\circ}^{T} c_{p_k}(T)\, dT}_{=h_s=\text{ sensible}} + \underbrace{\Delta h^\circ_{f,k}}_{\text{chemical}}, \tag{6.5}$$

where, for the k^{th} species, c_{p_k} is the isobaric heat capacity and $\Delta h^\circ_{f,k}$ is the formation enthalpy, and $\square^\circ$ represents the standard reference state. The mixture-averaged mass-based enthalpy is given by

$$h = \sum_{k=1}^{K} Y_k h_k \tag{6.6}$$

The second-to-last term on the RHS of Eq. (6.4) represents a pathway for heat flow even under isothermal conditions. This is due to the diffusion of species with different enthalpies. The final term in Eq. (6.4) is the rate of heat absorption by the

reaction, and can be calculated as

$$\dot{\omega}_T = \sum_{k=1}^{K} \dot{\omega}_k \Delta h^{\circ}_{f,k} \tag{6.7}$$

Although there are many different forms for the energy equations [105], Eq. (6.4) is preferred because its form is similar to that of non-reacting flow, with the addition of two additional terms to account for the reaction. This simplifies the procedure of introducing reactions into an existing non-reacting CFD solver.

6.2.2 Mass-Action Kinetics and the Arrhenius Rate Law

For a general set of elementary reversible reactions

$$\sum_{k=1}^{K} \nu'_{kj} X_k \underset{\mathrm{k}_{b,j}}{\overset{\mathrm{k}_{f,j}}{\rightleftharpoons}} \sum_{k=1}^{K} \nu''_{kj} X_k \quad \text{for } j = 1, \ldots, N_r, \tag{6.8}$$

which can, for example, look like

$$\mathrm{C_2H_3} + \mathrm{C_4H_8} \underset{\mathrm{k}_b}{\overset{\mathrm{k}_f}{\rightleftharpoons}} \mathrm{C_2H_4} + \mathrm{C_4H_7}, \tag{6.9}$$

the k^{th} species net species production rate is given by the law of mass action [90]

$$\dot{\omega}_k = \sum_{j=1}^{N_r} \dot{\omega}_{kj} = \frac{1}{\rho} \sum_{j=1}^{N_r} \nu_{kj} \left(\underbrace{k_{f,j} \prod_{k=1}^{K} [X_k]^{\nu'_{kj}}}_{\text{forwards}} - \underbrace{k_{b,j} \prod_{k=1}^{K} [X_k]^{\nu''_{kj}}}_{\text{backwards}} \right), \tag{6.10}$$

where N_r is the number of reactions involving species k, is the ν_{kj} are the net stoichiometric coefficient for the k^{th} species of the j^{th} reaction, $[X_k] = \frac{Y_k \rho}{M_k}$ is the molar concentration of species k, M_k is the molecular weight of species k, and $k_{f,j}$ & $k_{b,j}$ are the forward and backward reaction rates, described by a modified Arrhenius rate law

$$k_j = A_{f,j} T^{\beta_j} \exp\left(\frac{-E_{a,j}}{R_u T}\right), \tag{6.11}$$

where R_u is the universal gas constant, and for the j^{th} reaction, $A_{f,j}$ is the pre-exponential factor and $E_{a,j}$ is the activation energy.

The Jacobian Matrix

The Jacobian matrix is a critical component of stiff time-integration schemes [106] as well as in the construction of reduced chemical mechanisms. Most reacting flow codes use Newton's method to solve the implicit difference equation [107] for each species. This requires the Jacobian matrix to be constructed and inverted at each time step. It is defined as

$$\mathbf{J} = \frac{\partial \dot{\boldsymbol{\omega}}}{\partial \boldsymbol{Y}}, \tag{6.12}$$

where $\mathbf{J} \in \mathbb{R}^{K \times K}$, $\boldsymbol{Y} = [Y_1, \ldots, Y_K]^\top$ and $\dot{\boldsymbol{\omega}} = [\dot{\omega}_1, \ldots, \dot{\omega}_K]^\top$. The Jacobian ultimately represents the sensitivity of the species' production rates to perturbations in the thermochemical state $[T,\ p,\ Y_1, \ldots, Y_K]^\top$.

For the k^{th} species, the chemical time-scale and hence the Damköhler number is determined from the eigenvalues λ_k of the Jacobian matrix. The first Damköhler number is defined as

$$Da_{\text{I},k} = \frac{\tau_f}{\tau_{c_k}} = \tau_f |\mathfrak{Re}(\lambda_k)| \tag{6.13}$$

6.2.3 Physical Modelling Assumptions

In this subsection, the physical assumptions and approximations applied to the governing equations are outlined and justified. These are important as they also translate to the `ChemZIP` methodology introduced later in Chapter 7.

Thermodynamics

(A.1) *Local thermodynamic equilibrium:* It is assumed that within a macroscopic subvolume, the mixture is locally in thermodynamic equilibrium [108, 109]. Because the thermodynamic time scale is much faster than that of the chemistry, the chemical composition is assumed to be locally frozen from the perspective of the thermodynamics. This allows classical thermodynamic relations to be applied (*e.g.,* $R = c_p - c_v$) and also simplifies the calculation of thermophysical properties. For

example, although the isobaric heat capacity is formally defined as

$$c_p = \left(\frac{\partial h}{\partial T}\right)_p \equiv \underbrace{c_{p_f}}_{\text{frozen}} + \underbrace{\sum_{k=1}^{K} h_k \left(\frac{\partial Y_k}{\partial T}\right)_p}_{\text{chemical}}, \tag{6.14}$$

under local thermodynamic equilibrium assumptions, the last term is zero and the equation can be simplified to

$$c_p(T, Y_1, \ldots, Y_K) = \underbrace{c_{p_f}(T, Y_1, \ldots, Y_K)}_{\text{frozen } c_p} \equiv \sum_{k=1}^{K} Y_k c_{p_k}(T), \tag{6.15}$$

where $c_{p_k}(T)$ is the heat capacity of species k determined from NASA polynomials fits [104]. Similar assumptions can be made to define the frozen speed of sound, the frozen ratio of heat capacities, the frozen isochoric heat capacity, etc. [108]. This will significantly simplify adding chemistry with the `ChemZIP` methodology (see Chapter 7) to an existing non-reacting CFD solver and considerably increase performance.

(A.2) *Ideal gas equation of state:* Under the conditions of (A.1), the governing equations are closed with the ideal gas law combined with Dalton's law,

$$p = \rho R_u T \sum_{k=1}^{K} \frac{Y_k}{M_k} \tag{6.16}$$

The gas mixture is safely approximated as a semi-perfect superheated vapour. Non-ideal gas effects are negligible as a result of a high reduced temperature ($T_r = T/T_{\text{critical}} > 2.5$) and a low reduced pressure ($p_r = p/p_{\text{critical}} < 0.03$) in the chemical reaction zone for many processes (see Sec. 4.2).

Diffusion

(A.3) *Pure Fickian molecular mass diffusion approximation:* A complete description of multicomponent molecular mass diffusion is the Stefan-Maxwell equations [90, 110]. These equations account for the fact that each species diffuses relative to each

other species with a separate binary diffusion coefficient for each pair of species. To determine the binary diffusion coefficients, this results in a $(K-1)\times(K-1)$ linear system in each mesh cell. This needs to be solved, which is costly $\mathcal{O}(K^2)$, especially for a large number of species. Therefore, as an approximation, we assume each species diffuses relative to the *same* bulk gas mixture, so there is only one diffusion coefficient per species. As such, a Fickian molecular diffusive flux model $\rho\bar{D}_k\nabla Y_k$ (see Eq. (6.2)) can be used and a linear system no longer needs to be solved [110, 111]. This is a reasonable and common approximation because, apart from very close to the walls, precursor CFD indicates that the turbulent diffusivity overshadows the laminar diffusivity by two to three orders of magnitude, especially in the highly turbulent turbo-reactor environment.

(A.4) *Constant species Lewis number:* As before, since turbulent transport overwhelms molecular mass transport, a further simplifying assumption is to fix the Lewis number for each species [111, 112],

$$Le_k = \frac{\bar{\lambda}}{\rho c_p \bar{D}_k} \approx \text{constant} \tag{6.17}$$

The species' Lewis number can then be used to efficiently calculate the molecular mass diffusivity $\bar{D}_k$. While the constant-Lewis-number assumption is often suitable for combustion [105], for a steam cracking reaction, simple CANTERA simulations have indicated that $Le_k = Le_k(T, p, Y_1, \ldots, Y_K) \neq$ constant, and it can vary by 10 – 50 % throughout the reaction evolution for each species. Therefore, this approximation may need to be refined in future work.

(A.5) *Neglect cross-species heat flux:* The term

$$q_{st} = \nabla \cdot \left(\rho \sum_{k=1}^{K} h_k \bar{D}_k \nabla Y_k \right)$$

in Eq. (6.4) is neglected on the assumption that it is overwhelmed by the reaction heat absorption term $\dot{\omega}_T$ (see Eq. (6.7)) [105]. Although this is justified for combustion

due to very high heat release rates, for steam cracking $\dot{\omega}_T$ is lower and thus q_{st} may need to be included in future work.

Turbulence-Chemistry Interaction

(A.6) *Neglect turbulence-chemistry interaction:* Turbulence-chemistry interactions (TCI) come about from the non-linearity of the net species production rate in Eq. (6.2). Therefore,

$$\overline{\dot{\omega}_k}(T, p, Y_1, \ldots, Y_K) \neq \dot{\omega}_k \left(\overline{T}, \overline{p}, \overline{Y}_1, \ldots, \overline{Y}_K\right),$$

where the overbar $\overline{\square}$ indicates time/Reynolds averaging. For the RANS equations, a closure model is needed to relate the mean temperature $\overline{T}$ and the mean mass fractions to the Reynolds-averaged net species production rates. Due to the exponential temperature dependency of the rate constants (see Eq. (6.11)), turbulent temperature fluctuations are much more important than turbulent concentration or pressure fluctuations [113]. The reaction only responds to turbulent temperature fluctuations if the characteristic chemical time-scale is similar to the turbulent mixing time-scale (see Fig. 6.2).

Although these interactions can be important in many flows, a study by Rayniers *et al.* [113] suggests that, for steam cracking reactions in tubular coils, neglecting TCI only increased the yield by $0.1 - 0.2$ wt % absolute with turbulence intensity around 3% – 4%. Therefore, as a starting point, we neglect TCI. However, since the turbulence level and gradients in the turbo-reactor exceed those in a conventional coil by one to two orders of magnitude, this assumption may need to be refined in future work. For details on closure models available for turbulent reacting flows, the reader is referred to Fox [114].

In summary, these modelling approximations are useful because (1) they simplify the implementation of chemistry in an existing non-reacting CFD solver (2) they reduce computational cost and (3) they simplify the verification processes of `ChemZIP` against the FOM later. In future work, it may be necessary to relax some of these assumptions, especially, (A.4), (A.5) and (A.6).

6.3 Accelerated Aerochemical Modelling Literature Survey

Direct integration of detailed chemical kinetics is not a viable strategy to bring aerochemical CFD to the aerodynamic design optimisation stage. Therefore, "the mission" is to design an efficiently coupled aerochemistry modelling methodology that captures the rich dynamics and broad thermochemical state-space coverage provided by a comprehensive mechanism, while reducing computational cost relative to the FOM by two to three orders of magnitude.

Error-free acceleration strategies through improvements in numerical linear algebra, numerical methods and hardware (*e.g.,* GPUs, see [115, 116]) have been virtually exhausted for reacting flow modelling. Today, most state-of-the-art solvers use highly-efficient variable-order (from 1 to 5), variable-time-step multi-step methods (*e.g.,* CVODE [107]) to combat numerical stiffness. These methods require evaluating the Jacobian matrix. Its evaluation has been accelerated by automatically generating analytic formulations instead of using a finite-difference approximation, gaining a performance boost of a factor of 3 to 7 [106]. A particularly fruitful acceleration strategy came from McNenly *et al.* and [117] Walker *et al.* [118] who used adaptive preconditioning to increase the sparsity of the Jacobian matrix. This allowed specialised sparse non-linear equation solvers (*e.g.,* GMRES) to be used to improve performance. It was demonstrated that orders of magnitude speed-ups were possible for large mechanisms [118], but this does not alleviate the cost of hundreds of additional transport equations for CFD.

6.3.1 *Physics-Driven* Strategies

The proceeding literature review is broken down into physics-based acceleration strategies (this subsection) and machine-learning-based techniques (in the next subsection Sec. 6.3.2). It should be noted that many, if not all, of the methods below are often combined to achieve optimal speed-up. In general, it is difficult to rigorously compare speed-up factors because they depend significantly on the mechanism size.

Global Mechanisms

At the most reductive end of the spectrum of acceleration strategies, there are global reaction mechanisms [119]. The reaction mechanism is simplified to 5 – 6 species & 2 – 4 reactions, and their rate constants are fitted to experimental data. While this certainly makes reacting flow computationally tractable, it is well known that so much of the rich functional expressiveness is lost that global mechanisms cannot accurately predict even the most basic reaction properties [112].

Skeletal Mechanism Reduction and Dynamic Adaptive Chemistry

A less restrictive approach is to construct a skeletal mechanism. A popular strategy to automatically and systematically build reduced mechanisms is a set of graph-based techniques developed by Lu & Law [120] in 2004 and Pepiot & Pitsch [121] in 2008. An open source implementation called PyMars can be found at [122]. The reaction network is formulated as a directed relation graph, with interaction coefficients (edges) between each species (nodes) defined by normalised reaction rates. Inconsequential species and chemical paths are filtered if their coupling strength with a selected number of target (major) species is below a user-defined threshold. Figure 6.5 shows that a remarkably good prediction of the ignition delay can be obtained for 850-species iso-octane combustion with a quarter of the number of species contained in the detailed mechanism.

However, skeletal mechanisms can still be relatively large. For example, Lu & Law [123] showed that an n-heptane mechanism with 561 species could only be reduced to 188 species with an error of 20%. This remains a significant computational burden for 3-D viscous turbulent reacting flow simulations. Moreover, skeletal mechanisms are often only valid in restricted regions of the thermochemical space on which they were developed. This makes them inadequate for broad parameter space exploration studies. Specifically in the turbo-reactor, the complexity of the aerothermal flowfield means that more sub-mechanisms within the reaction network are active than for simpler flows, further impeding the effectiveness of skeletal mechanisms. This challenge can be cleverly addressed by building skeletal mechanisms "on-the-fly", locally within each cell of the mesh using the dynamic adaptive chemistry method (DACM) from Liang, Stevens &

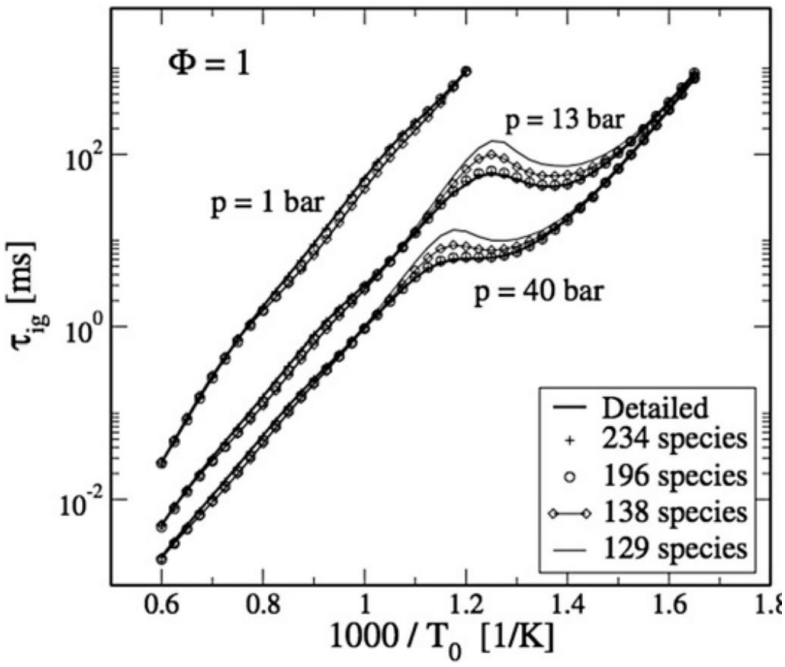


Figure 6.5: Comparison of ignition times obtained with a detailed iso-octane mechanism (850 species) and skeletal mechanisms of various sizes (reproduced from [121])

Farrell [124]. Although this method brings an additional computational overhead, a 30× performance benefit is still possible for a mechanism with 588 species [124].

In situ Adaptive Tabulation

An effective method to avoid evaluating the reaction network (see Eqs. (6.8)–(6.11) in Sec. 6.2.2) at each time integration step is the tabulation of the net species production rates $\dot{\omega}_k$ and its Jacobian matrix $\mathbf{J}$ as a function of the thermochemical state. Offline tabulation of the entire realisable composition phase space is impossible from a memory and retrieval perspective. Therefore, a pivotal method proposed by Pope in 1997 [125] was *in situ* adaptive tabulation (ISAT). This was later improved by Lu & Pope in 2009 [126]. In this approach, within each mesh cell, a lookup table is grown in real-time as the simulation progresses so that it only covers the relevant region of the composition space. This is possible in part because, as will be shown later, the chemistry lies in a low-dimensional region of the composition space region, *i.e.,* a manifold.

An effective tabulation scheme must be designed to maintain interpolation errors below a user-defined threshold while minimising the size of the lookup table (LUT).

Therefore, a key component of ISAT is estimating the error and then controlling it within specified bounds. New entries are only added to the table if errors are above a specified threshold *and* the tracked *region of accuracy* cannot be appropriately modified to contain the new query point.

Although ISAT has been successful in achieving two to three orders of magnitude speed-ups relative to direct integration, it can only be used with small mechanisms of less than 50 species. This is because the storage and retrieval scales quadratically with the number of species [127]. Furthermore, its performance degrades for highly unsteady flows—which are characteristic of flows in the turbo-reactor—where pre-tabulated entries may not be able to be re-used [128].

Dynamic Cell Zoning Method

In the context of 2D/3D reactive flows, Liang, Stevens & Farrell [129] introduced the dynamic cell zoning (DCZ) method in 2009 in which mesh cells with similar thermochemical conditions are grouped and the chemical kinetics evaluated only once. This can reduce the number of evaluations of the reaction network by a factor proportional to the number of mesh cells evaluated as one. Reyniers *et al.* [113] showed that a speed-up of a factor of 50 to 190 is possible (for a mechanism with 44 species) while keeping the relative error for feed conversion below 0.05%. The DCZ method works well for flows with more homogeneous conditions (*e.g.,* tubular pipe flow), but less so when for cases with a large degree of non-uniformity like the turbo-reactor.

Intrinsic Low-Dimensional Manifold

Intrinsic low-dimensional manifolds (ILDM) were introduced in 1992 by Maas & Pope [130] and refined by Ren & Pope [131] and Bongers, Van Oijen & De Goey [132] to address the challenge of tabulating chemical kinetics when the composition space $(Y_1, \ldots, Y_K)$ is high-dimensional (K can be on the order of 1000 species). Using the ILDM approach, the K-dimensional composition space is projected onto a reduced K_m-dimensional manifold space, where $K_m \ll K$. This is achieved by leveraging ideas from mathematical modelling of dynamical systems and time-scale analysis. Figure 6.6 illustrates that for various initial

conditions, fast time-scales are quickly attracted to a low-dimensional manifold in phase space governed by the trajectories of slow time-scales. Therefore, by performing a singular value decomposition of the Jacobian (see Eq. (6.12)), a new composition space can be constructed based only on a limited number of slow modes. It is assumed that the fast modes are in quasi-steady state, allowing the differential equation to be converted into an algebraic one. The manifold is then tabulated as a pre-processing step.

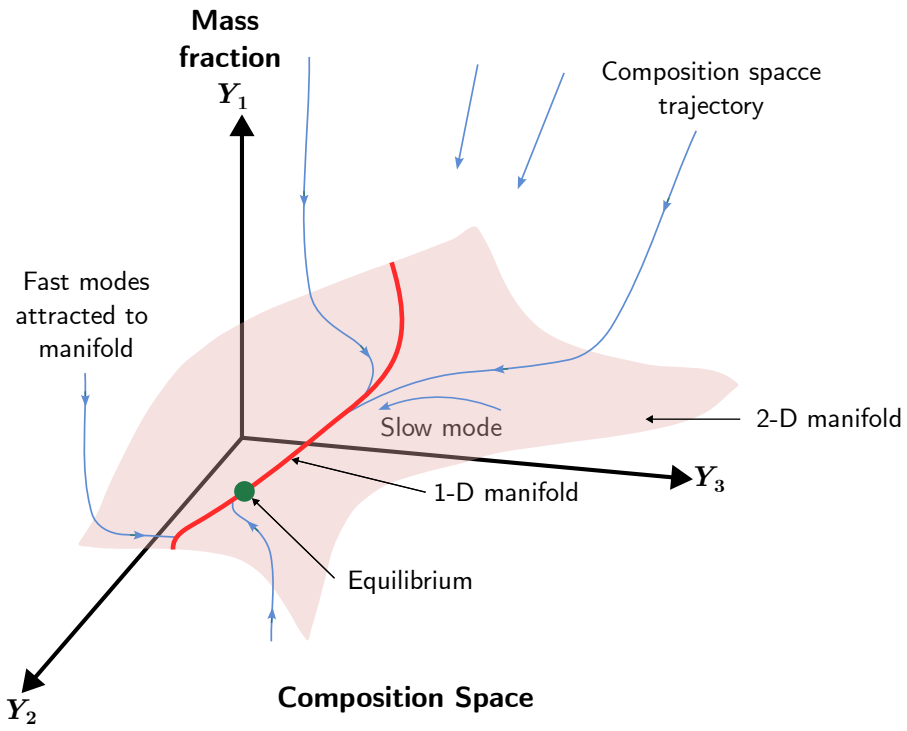


Figure 6.6: Schematic illustration of a low-dimensional manifold in composition space

Flamelet Generated Manifold

The original ILDM method did not account for convection and diffusion for the short time-scale modes (which were assumed to be in quasi-steady state). This led to errors in predicting the colder regions of laminar flames where the reaction no longer dominated transport. Therefore, ILDM was superseded by the flamelet-generated manifold (FGM) approach proposed by Van Oijen in 1999 [133] and explained in detail in Ref. [134]. Today, this has become one of the most popular approaches for flame calculations [133]. The FGM methodology combines the ideas of low-dimensional manifolds from ILDM with 1-D *flamelets*. Assuming the reaction zone of a multidimensional flame is thin,

it can be approximated as an ensemble of 1-D flamelets. Therefore, the trajectory in composition space is assumed to be close to that of 1-D flames.

Typically, two control variables are used to track progress through a library of offline-tabulated 1-D flamelet calculations. A common set of tracking variables are the mixture fraction (which ranges from 0 in the unburnt state to 1 in the fully burnt state) and the reaction progress variable. Their URANS transport equations can be described as follows:

$$\frac{\partial(\rho\mathcal{Z})}{\partial t} + \nabla\cdot(\rho\vec{\mathbf{u}}\mathcal{Z}) = \nabla\cdot\left(\left(\rho\bar{D}_{\mathcal{Z}} + \frac{\mu_t}{Sc_t}\right)\nabla\mathcal{Z}\right) \tag{6.18}$$

$$\frac{\partial(\rho\mathcal{Y})}{\partial t} + \nabla\cdot(\rho\vec{\mathbf{u}}\mathcal{Y}) = \nabla\cdot\left(\left(\rho\bar{D}_{\mathcal{Y}} + \frac{\mu_t}{Sc_t}\right)\nabla\mathcal{Y}\right) + \dot{\omega}_{\mathcal{Y}}, \tag{6.19}$$

where one possible definition for the mixture fraction and the progress variable are

$$\mathcal{Z} = \frac{sY_{\text{Fuel}} - Y_{\text{Oxi.}} + Y^0_{\text{Oxi.}}}{Y^0_{\text{Fuel}} + Y^0_{\text{Oxi.}}} \tag{6.20}$$

and

$$\mathcal{Y} = \frac{Y_{\text{H}_2\text{O}}/M_{\text{H}_2} + Y_{\text{H}_2\text{O}}/M_{\text{H}_2\text{O}} + Y_{\text{CO}_2}/M_{\text{CO}_2}}{Y_{\text{H}_2\text{O},eq}/M_{\text{H}_2} + Y_{\text{H}_2\text{O},eq}/M_{\text{H}_2\text{O}} + Y_{\text{CO}_2,eq}/M_{\text{CO}_2}} \tag{6.21}$$

respectively [135]. Given a pair of $\mathcal{Z}$ and $\mathcal{Y}$, it is possible to extract all species mass fractions and net production rates from the pre-tabulated look-up table. Figure 6.7 demonstrates that FGM can provide highly accurate predictions of laminar flame chemistry. The FGM methodology has been extended to turbulent flames by adding two additional transport equations to account for TCI, as demonstrated by Klapdor [135] for a combustor-turbine interaction test case.

Flamelet-generated manifolds provide a powerful computational cost minimisation, reducing the number of transport equations from K to 2 (for laminar flames) while also eliminating the cost of evaluating the reaction network using a LUT. However, in its current form, FGM is only applicable to combustion and therefore cannot be used for modelling chemistry in the turbo-reactor. Despite this, the mathematical and physical principles used in both FGM and ILDM will be drawn upon later in the development of `ChemZIP` (Chapter 7), generalising these methodologies using machine learning. A particularly

useful idea is to track progress through a compressed chemical kinetic manifold using a limited number of transported quantities.

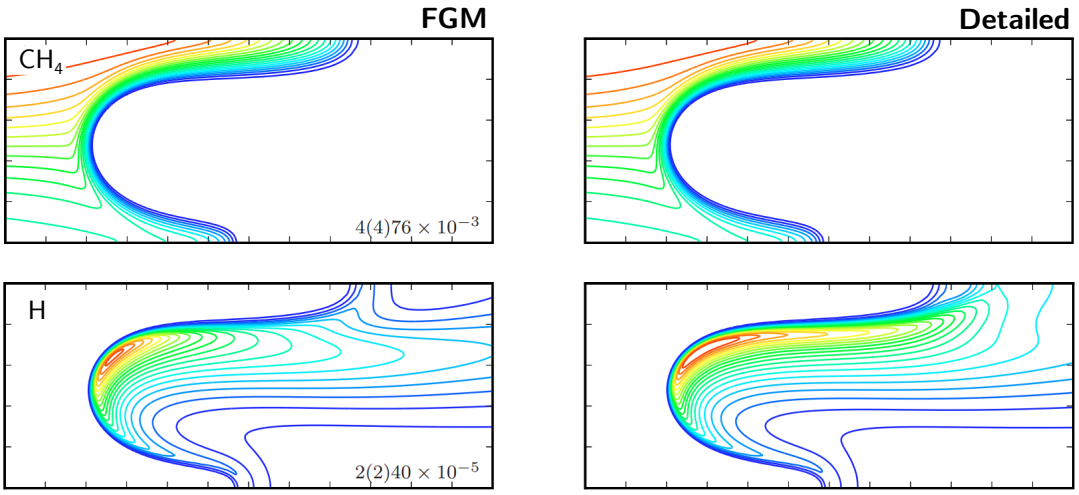


Figure 6.7: Verification of FGM against detailed chemistry for a 2D premixed laminar methane flame (reproduced from [134])

6.3.2 *Data-Driven* Strategies

Physics-driven approaches have provided substantial computational cost minimisation while maintaining good levels of accuracy. In the search for greater acceleration capabilities, data-driven approaches have been explored by many researchers. Some of these approaches are reviewed below as well as summarised in Ref. [136].

Artificial Neural Networks

In the early 2000s, the memory footprint of a LUT was a key limiting factor. This motivated the use of function approximators. It is well known that the number of coefficients for polynomial regression scales exponentially with dimensionality [137]. Artificial neural networks (ANN) bypass this so-called *curse of dimensionality* [137] and can approximate the functional mapping

$$\dot{\boldsymbol{\omega}} \approx \mathscr{F}_{\mathrm{NN}}(T, p, \boldsymbol{Y}), \tag{6.22}$$

for the reaction network with many fewer parameters. By storing information in only a small number of parameters, neural networks (NN) were able to reduce memory

consumption 100-fold compared to ISAT for a similar level of accuracy [138]. However, today, memory bottlenecks are less of a concern. The main focus is instead on the CPU time for evaluating/retrieving the net species production rates. There are mixed reports on the speed of ANNs compared to ISAT [139, 140]. The outcome depends heavily on accuracy requirements and the number of species in the mechanism. To address this confusion, Nikolaou, Vervisch & Domingo [141] developed a mathematical framework to identify a rigorous selection criterion.

The first successful implementation of ANNs to accelerate chemical kinetics dates back to 1999 from Christo *et al.* [142]. Since then, there have been numerous implementations of ANNs, particularly in combustion [139, 143–146]. The output of the neural network is either the net species production rates (see Eq. (6.22) and Fig. 6.8(a)) or the change in chemical state to the next time step (see Fig. 6.8(b));

$$d\boldsymbol{Y} \equiv \boldsymbol{Y}(t+\Delta t) - \boldsymbol{Y}(t) = \int_t^{\Delta t} \dot{\boldsymbol{\omega}}\, dt \approx \mathscr{F}_{\mathrm{NN}}(T, p, \boldsymbol{Y}) \tag{6.23}$$

This requires operator splitting [127] to decouple the time integrator used for the convection/diffusion operator from that used for the chemical source terms $\dot{\boldsymbol{\omega}}$ (see Eq. (6.23) above). A simplified formulation of operator splitting is outlined as follows:

$$\frac{d\boldsymbol{Y}^*}{dt} = \boldsymbol{H}(\boldsymbol{Y}^*, T, p), \quad \boldsymbol{Y}^*(t_n) = \boldsymbol{Y}(t_n), \quad t \in [t_n, t_n + \Delta t], \tag{6.24}$$

$$\frac{d\boldsymbol{Y}^{**}}{dt} = \dot{\boldsymbol{\omega}}(\boldsymbol{Y}^{**}, T, p), \quad \boldsymbol{Y}^{**}(t_n) = \boldsymbol{Y}^*(t_{n+1}), \quad t \in [t_n, t_n + \Delta t], \tag{6.25}$$

where $\boldsymbol{H}$ is the sum of the convection and diffusion operators. Equation (6.25) is the equation for a 0-D homogeneous, constant-pressure reactor, solved in each mesh cell.

Training data for ANNs is typically generated from either 1-D flame and/or 0-D reactor calculations using a detailed kinetic model. The thermochemical space is sampled by varying the initial equivalence ratio, temperature and pressure, and the dynamic response is stored in a library with order 1 to 10 million samples. It is challenging to sufficiently cover the thermochemical space, especially since the species concentrations

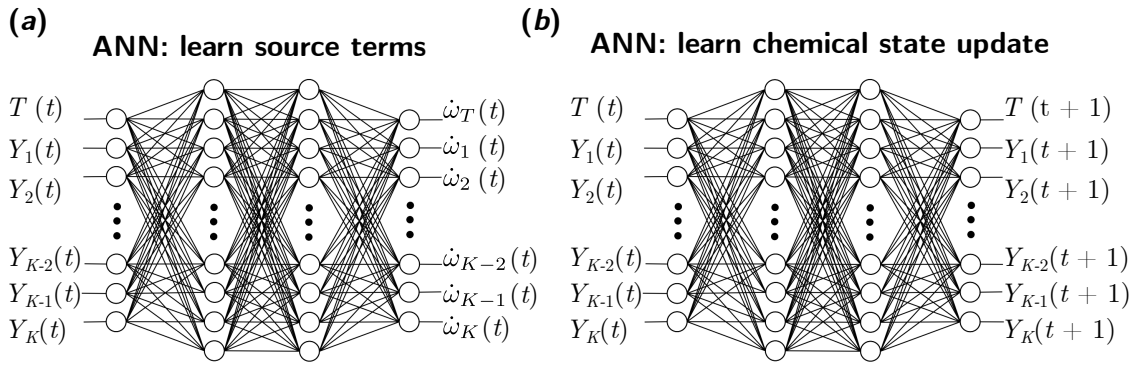


Figure 6.8: Artificial neural network schematic illustration

are highly multi-scale. Therefore, specialised sampling methods have been developed to sample uniformly across many orders of magnitude [147, 148].

For reacting flow simulations, the neural network has to be evaluated at each time-step and in each mesh cell. This inference time scales with the number of NN parameters (*i.e.,* trainable coefficients), which can be $\mathcal{O}(5,000 - 1{,}000{,}000)$ to accurately approximate the entire thermochemical space. Therefore, many ANN implementations partition the thermochemical phase space into sub-regions. For each sub-region, a smaller neural network is trained [149, 150]. Because it only has to approximate a small region of the thermochemical space, fewer NN parameters can be used, reducing the inference time during a reacting flow simulation.

A common strategy to bound the NN solution is to encourage elemental mass conservation during optimisation of the NN parameters [151]. Introducing physics-based prior knowledge regularises the network and improves generalisability; however, these additional physical constraints can limit the functional expressiveness of the network and lead to reduced accuracy.

A key limitation of the "vanilla" ANN is that it does not alleviate any of the numerical stiffness (*i.e.,* $S = \tau_{\text{max.}}/\tau_{\text{min.}}$) and therefore provides little to no performance benefit over ISAT. This is addressed below.

Neural Operators

Recently in 2024, Goswami *et al.* [152] introduced an elegant approach based on deep neural operators [153], which can alleviate the stiffness of kinetics. To achieve this,

time-step information is encoded into the neural network architecture, allowing the solution operator/propagator itself to be learnt for arbitrary time instances rather than just the chemical source terms. The action of the ODE solution propagator, denoted as $\mathscr{F}_{\mathrm{NN},t}^{t+\Delta t} : \mathbb{R}^K \to \mathbb{R}^K$, is defined as follows:

$$\boldsymbol{Y}(t+\Delta t) = \mathscr{F}_{\mathrm{NN},t}^{t+\Delta t}(\boldsymbol{Y}(t)) \tag{6.26}$$

Note that this is not the same as Eq. (6.23), which is for a fixed time-step size. Therefore, the trained solution operator can naturally interpolate and "jump ahead" to much larger time steps, similar in size to the fluid time step, despite being trained on much finer temporal resolution data required for the accuracy and stability of conventional ODE solvers. Goswami demonstrated that highly accurate predictions are possible with a time step 1000 times larger than that needed for traditional stiff integrators. This makes data-driven approaches more favourable than the physics-based approaches discussed in Sec. 6.3.1—and their attractiveness will be further enhanced with data-driven dimensionality reduction introduced below.

Principal Component Analysis and Autoencoders

The data-driven approaches mentioned above do not mitigate the cost of transporting a large number of species mass fractions (which can be $\mathcal{O}(1000)$ for a detailed mechanism). Furthermore, very large neural networks are required to approximate the high-dimensional reaction dynamics, leading to long-latency inference times. Therefore, it is still prohibitive to incorporate detailed chemistry into high-fidelity fluid solvers.

To address this, Sutherland & Parente [154] used principal component analysis (PCA) to learn a low-dimensional natural coordinate system (or equivalently a manifold) to project the composition space onto. This is similar to ILDM and FGM, but more general and automated. The reduced dimensionality state vector ($M \ll K$) is formed as a weighted linear combination of mass fractions from the original composition space, exploiting the fact that many species are correlated. Many works have successfully demonstrated combining PCA with non-linear regression techniques, such as ANNs and

neural operators, to approximate the dynamics in the low-dimensional composition space [155–158]. However, poor manifold topology has been a common challenge, which can lead to overlapping regions of the manifold, and thus non-uniqueness. This means that, given a state, the dependent variable can take multiple values. To address this, Zdybał, Parente & Sutherland [158] recently developed an elegant cost function to evaluate and ultimately improve low-dimensional manifold quality/topology, as shown in Fig. 6.9.

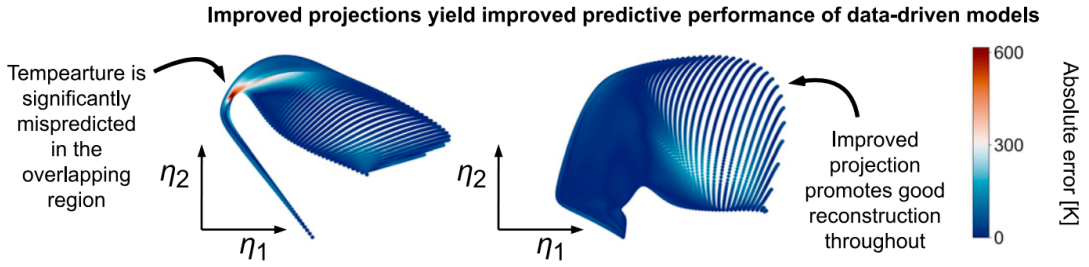


Figure 6.9: Improving low-dimensional projected manifold quality (reproduced from [158])

Autoencoder NNs have also been used to learn low-dimensional *non-linear* representations of the composition space [152, 159–161], compared to the *linear* representations of PCA. This enables more accurate compression and greater stiffness mitigation than PCA. Grassi *et al.* [162] combined autoencoders (AE) with a simplified Arrhenius-based chemical mechanism in latent space. This involved learning the Arrhenius rate coefficients for an artificial "meta" reaction A $\rightleftharpoons$ B in latent space. While this provides an interpretable low-dimensional manifold, it restricts functional expressiveness in latent space, and can severely limit accuracy. This approach does not work for complex systems such as the turbo-reactor.

These methodologies have led to sufficiently powerful computational cost minimisation that high-fidelity combustion simulations have been performed under engine-relevant conditions [163, 164].

6.4 Summary

This chapter has highlighted the importance of capturing aerochemical interactions early at the design level without sacrificing accuracy, and introduced the basic equations of reacting flow and key assumptions that will also translate to `ChemZIP`.

Direct integration of the fully coupled reacting flow system is intractable—particularly for high-carbon-number feeds for steam cracking in which the kinetic models are large and multi-scale. Therefore, various physics-based approaches have been introduced to accelerate aerochemical simulations, some of which are default options in commercial solvers like FLUENT. This limitations of traditional acceleration approaches have been highlighted: skeletal mechanisms are too large and DACM incurs overheads from on-the-fly mechanism reduction; ISAT is ineffective for a reasonable number of species and its performance degrades for transient flows; DCZ is still more than 10 times slower than the levels of computational cost reduction needed; and FGM approaches are efficient, but only work for combustion modelling. In the turbo-reactor, the notion of a flamelet does not exist, and the mixture fraction and reaction progress variables are not defined.

Therefore, with the step-change advances in machine learning and its proliferation across computational science, the research focus has shifted to PCA/AE—ANN/neural-operator techniques in the quest for greater levels of acceleration and greater universality. Many researchers have demonstrated superior acceleration with data-driven approaches; for example, by combining PCA with neural operators, speed-ups of a factor of 2000 are possible for a 350-species mechanism [160].

In the next chapter, we will combine several of the data-driven approaches discussed above to develop a new methodology for efficiently coupling aerochemistry into the design cycle in the context of a new application for reactive flows: conducting chemical reactions within turbomachines. This is an exciting new application for reacting flow modelling, but it poses a new set of challenges.

Chapter 7

Efficiently Coupled Aerochemical Modelling Methodology

7.1 Why Machine Learning?

To enable aerochemical modelling for the first time at the design level, this work develops a multi-fidelity machine-learning-assisted efficiently coupled aerochemical modelling methodology (*i.e.,* `ChemZIP`). This leverages many of the techniques from the literature (see Chapter 6), but with some unique challenges and differences due to the novelty of the application. To date, acceleration methods have been developed mainly for combustion modelling. However, aerochemical interactions in combustion flows are considerably different from those along the bladed path of a turbomachine. In addition, while other approaches focus on accuracy, we focus on *low-latency deployment/inference* so that design optimisation is feasible, even if it means sacrificing accuracy while still capturing important trends.

A relevant question is: why choose machine-learning-based acceleration techniques (PCA, AEs, ANNs, etc.) over physics-based approaches (ISAT, DACM, skeletal mechanisms, etc.)? The main reason is, of course, greater computational acceleration. Generally, across the field of computational science, it is becoming clear that a key role of ML is developing fast approximate surrogate models for problems that involve tens of thousands of prohibitively costly simulation runs, but can tolerate some accuracy loss if trends are well predicted (*e.g.,* optimisation). In these settings, machine learning can be used to strike a good balance between accuracy and speed.

However, many engineers are reluctant to embrace ML because of its lack of interpretability, long training time-frames, and/or the absence of a deterministic vehicle for error control. While these are valid concerns, the discourse surrounding ML sometimes distracts practitioners from the fact that, when provided with a well-crafted and extensive

training database, ML provides a powerful means of information compression and an elegant bridge between different fidelity methods and physical models. In the context of this work, ML allows the dynamic response of detailed and accurate (but problematically high-dimensional) 1-D chemistry to be efficiently stored, allowing it to be accessed extremely fast by a high-fidelity aerodynamic solver. Compared with tabulation, NNs clearly provide a more efficient means of storing, compressing and smoothly interpolating between large amounts of high-dimensional information using only a limited number of parameters. This allows for much greater coverage of the thermochemical space without having to resort to *on-the-fly* methods that incur additional overheads. This is a critical advantage.

In the previous chapter, skeletal mechanisms, ILDM, and FGM were introduced as different but related routes to eliminate fast-reacting species that do not significantly influence the important target species (*e.g.,* ethylene). Effectively, a low-dimensional coordinate system is uncovered that accurately reproduces the thermochemical dynamics of the important target species. Data-driven compression (*e.g.,* through an AE) is ultimately an extension of these methods that can be applied more universally (compared to FGM), automatically, and while achieving greater levels of compression through non-linearity.

In an attempt to demystify *data-driven* compression, it is useful to draw parallels with the relevant and related *physics-based* acceleration approach familiar to the turbomachinery community: coupling unsteady thermofluid modelling with solid conduction *i.e.,* unsteady conjugate heat transfer (CHT). He & Oldfield [165] and He [166] demonstrated that by re-framing the CHT problem in the frequency (rather than time) domain, the fluid-solid coupling could be accelerated. This is possible because, from the perspective of the solid, a Fourier basis for wall temperature unsteadiness can be reliably compressed by discarding unimportant high-frequency harmonics. A similar argument applies to image compression where, again, the Fourier space is a naturally sparse coordinate system for visual data.

Data-driven approximation theory can achieve the same outcome. However, instead of having a predetermined and fixed set of basis functions (*e.g.,* a Fourier basis), a tailored coordinate system (and set of basis functions) that best sparsifies the given data is learnt during training. Often, there is no physical intuition to inform the selection of the

basis functions for a given problem (*e.g.,* chemical kinetics). This is where data-driven approaches provide an effective solution.

In this chapter, the implementation details of the new `ChemZIP` methodology will be described, including training data generation, the neural network architecture, and the kinetics-fluid coupling treatment. In the new `ChemZIP` methodology a NN architecture compresses the high-dimensional thermochemical space $\mathcal{O}(200-1000)$ onto a substantially lower-dimensional manifold $\mathcal{O}(3)$, embedded in the original state space, where the system dynamics can be tracked. In latent space, a reduced set of *meta-species* mass fractions are mapped to their corresponding source terms, which correspond to a projection of the physical dynamic response into latent space. This approach not only minimises the cost of evaluating the reaction network but also alleviates numerical stiffness and reduces the number of transport equations. Since this work is a proof-of-concept investigation, simpler implementation approaches have been adopted to accelerate the model development timeline. However, this will be refined in future work. For an introduction to modern ML concepts and NNs, Goodfellow, Bengio & Courville [137] and Bishop & Bishop [167] are recommended.

7.2 Training Data Generation

The quality of any ML method usually boils down to the quality and size of the suite training data. In computational science, training data generation is often the bottleneck of an ML pipeline. The `ChemZIP` methodology addresses this challenge by simplifying the chemistry model down to a perfectly-mixed, constant-area 1-D PFR while maintaining an accurate and detailed kinetic model with $\mathcal{O}(200-1000)$ species. By reducing the cost of a simulation, tens of thousands of different conditions for the detailed kinetic model can be predicted in just a few hours, thus eliminating the training data generation bottleneck.

Neural networks are ultimately just universal input-output function approximators. Therefore, even in a simplified 1D environment, if we generate training data in such a way that all possible realizable combinations of the *inputs* Y_k, T & p are mapped to the dynamic response (*outputs*) of the reaction network in terms of the net species production rates $\dot{\omega}_k$, as well as $\dot{\omega}_T$, c_p, R, & the molecular viscosity $\bar{\mu}$, then in theory, with a sufficiently

large NN function approximator, any thermochemical input state could be accurately mapped to the output reaction response. However, this is a challenge for two reasons.

First, when there are more than 50 species, this is a combinatorially hard problem. Second, the NN would be large and excessive inference latency would negate the original intent of the ML approach (*i.e.,* computational acceleration). Therefore, a heuristic approach is taken, and the phase-space exploration is constrained to turbo-reactor-relevant conditions. However, generating these conditions is a challenging problem. Figure 7.1 shows that the flow in the turbo-reactor experiences strong temperature gradients and therefore much more severe and diverse reaction response regimes compared to a conventional tubular coil in a heated furnace. Therefore, to ensure that the training database includes the range of complex dynamic response regimes, first, precursor non-reacting CFD is conducted to determine the envelope of realistic temperature gradients seen by a fluid parcel travelling along the bladed flow path (*e.g.,* across a shock). A numerical simulation showing realistic temperature gradients is shown in Fig. 7.3. This forms the first component of the overall training data generation workflow shown in Fig. 7.2.

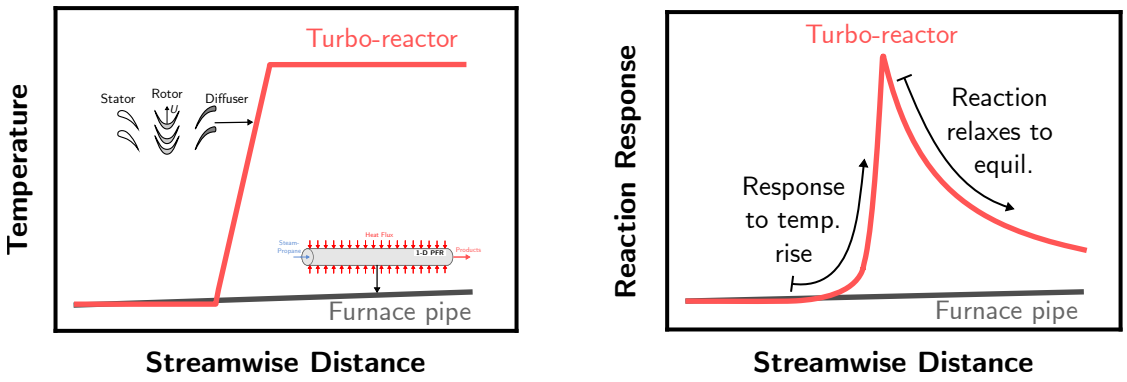


Figure 7.1: Schematic illustration of the reaction response to representative temperature gradients in the turbo-reactor compared to those in a tubular pipe of a conventional furnace

The next step is to randomly sample several thousand heat flux profiles, while remaining with the envelope of temperature gradients set by the precursor non-reacting CFD (see Fig. 7.3). By perturbing the reaction through sharp gradients in heat input to the reaction, the dynamic response of the reaction system ($\dot{\omega}_k$, $\dot{\omega}_T$, c_p, etc.) can be explored and recorded over a wide range of thermochemical conditions (see Fig. 7.2).

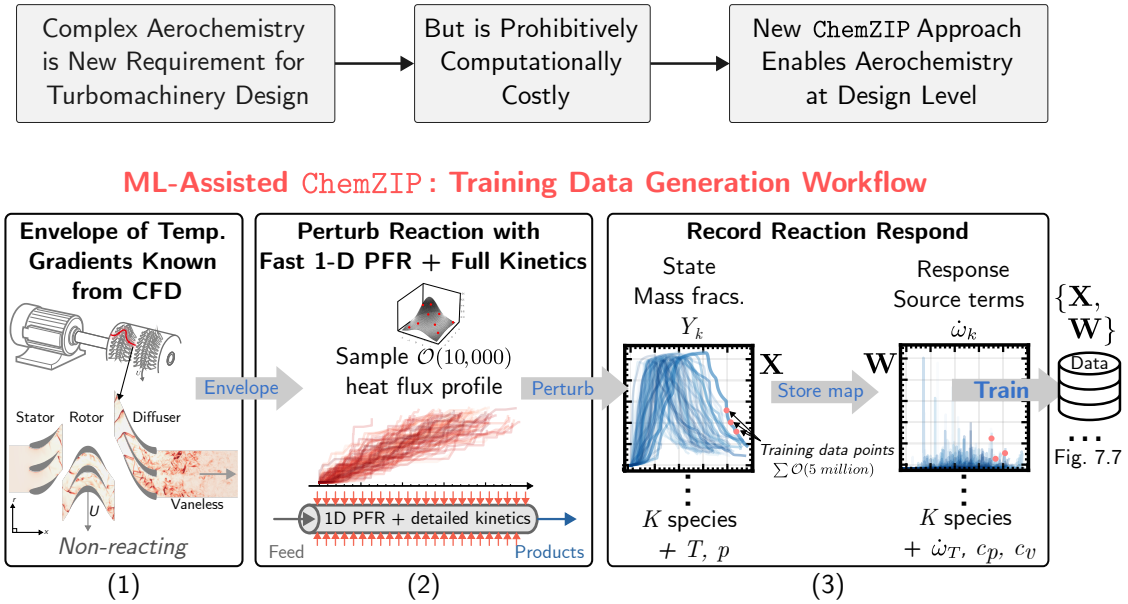


Figure 7.2: Schematic of the training data generation workflow for the ChemZIP methodology

These perturbations are not discontinuous but rather smooth changes in heat input resolved by a manageable number of integration points.

By repeatedly perturbing the reaction at different rates and at different times, the sensitivities (transfer functions) of the reaction system are generated; this should allow the downstream learning algorithm to construct a manifold that approximates the high-dimensional trajectory through composition space (see Fig. 6.6). To the best of our knowledge, this perturbation-based training data generation strategy has not yet been leveraged in the literature.

7.2.1 Sample Space Generation

The shape of each heat flux profile is set by N_{control} control points (see Fig. 7.4(a)), where N_{control} is typically around 10 – 20. The same methodology for defining the heat flux profile defined in Sec. 4.1.1 and illustrated in Fig. 4.1 is also adopted here. At each independent control point, the following control parameters can be varied to shape the heat flux profile: the residence time between control points (Δt), the heat input level (Δh_0), the heat input ramp rate, and the heat loss level (see Figs. 4.1 & 7.4). The heat input can be either positive or negative to simulate both heating and cooling. A typical envelope for these control parameters is shown in Table 7.1. In addition, the

pressure level can also float. If $N_{\text{control}} = 15$, this means that the total dimensionality of the sample space would be $15 \times 5 = 75$-D.

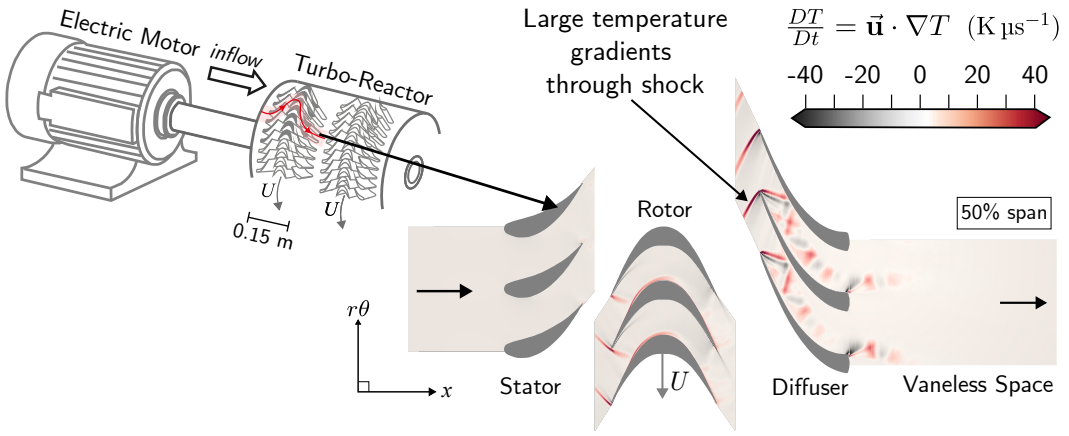


Figure 7.3: Typical temperature gradients $\frac{DT}{Dt} = \vec{\mathbf{u}} \cdot \nabla T$ (K µs^{-1}) seen by a fluid parcel travelling along the turbo-reactor gas path computed using non-reacting CFD

Table 7.1: Sample space parameter ranges for each heat flux profile control point

Residence time	Heat inp. level	Heat inp. rate†	Heat loss level
0.001 – 15 ms	−50 – 650 kJ kg^{-1}	0.001 – 100%	−30 – 30 kJ kg^{-1}

† this is set as a percentage of the heat input level

Sobol sampling [168] is used to distribute the samples uniformly in this high-dimensional (*e.g.,* 75-D) space. On the order of 5,000 – 10,000 heat flux profiles are randomly generated. To diversify the sample space, 2,500 samples are added where the heat flux is constant and only its magnitude is randomly varied on a log scale. For control parameters that are allowed to change over many orders of magnitude, such as the "residence time between control points" (see Table 7.1), sampling can be performed on a logarithmic (rather than linear) scale to spread the samples more effectively [147].

7.2.2 Constant-Pressure 1-D PFR Simulations

For each heat flux profile realisation, a constant-pressure 1-D PFR simulation is run (see Figs. 7.2 & 7.4) using the same method described earlier in Sec. 4.1.1 and illustrated in Fig. 4.1. The method described is agnostic to the kinetic model and/or initial composition used. In the current workflow, RMG-PY is used to generate detailed kinetic models. For development purposes, a 65-species propane kinetic model is used.

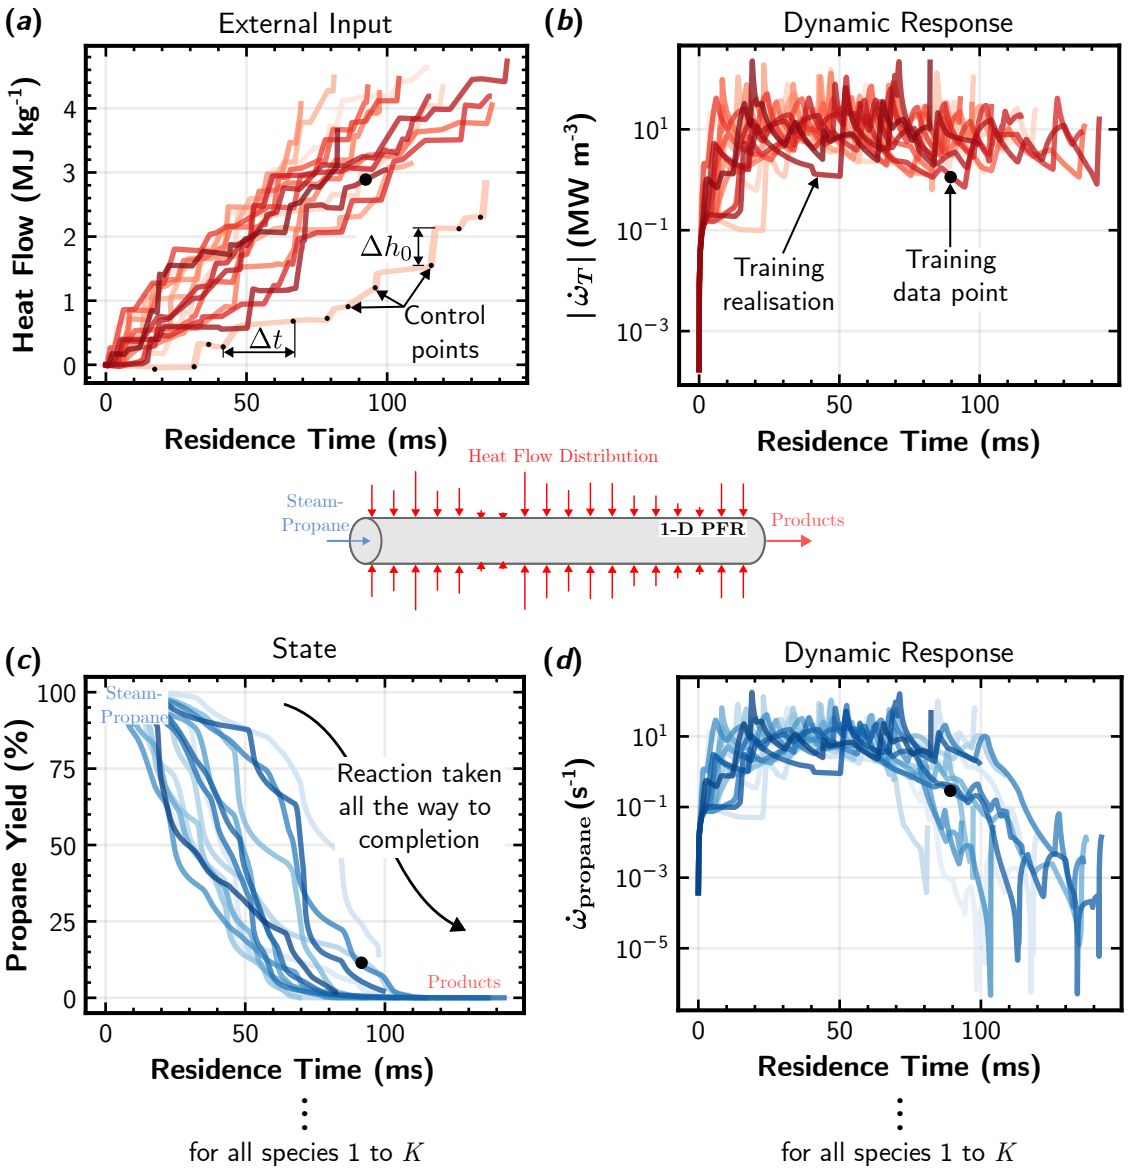


Figure 7.4: A selection of 10 (out of 10,000) example training data 1-D PFR reactor simulations showing **(*a*)** the heat flow **(*b*)** the reaction heat absorption rate on a log scale **(*c*)** the dry propane yield and **(*d*)** the propane net production rate on a log scale

The training data generation module of `ChemZIP` has been parallelised using `MPI4Py`. Although parallelisation is simple because each 1-D PFR simulation is independent, memory bandwidth is a bottleneck for large mechanisms. However, for mechanisms with fewer than a few hundred species, the training database can be populated within a few hours on an 18-core Intel Cascade Lake CPU.

Currently, the initial temperature ($T_{\text{initial}} = 650\,^{\circ}\text{C}$, which is below reaction activation), initial feed composition (100% propane) and steam ratio (SR = 0.5) are fixed, but can

vary if necessary. For each realisation, the simulation is shutdown early if the temperature exceeds a user-defined limit of $T_{\text{max.}} = 1150\,^{\circ}\text{C}$, which is 15% higher than the maximum allowable operating temperature of the turbo-reactor.

For each heat flux profile, $\mathcal{O}(1000)$ time steps are taken and the thermochemical state and dynamic response are stored in a training database. This results in a large corpus of training data, where $N = \mathcal{O}(5 \text{ million})$ samples for each variable. To allow a larger number of training data points to be clustered around regions of more complex reaction dynamics (which are thus challenging to capture), CANTERA is set to take adaptive time steps. The underlying time integrator naturally distributes more time steps in regions of high gradients where the truncation error rises above a threshold, as illustrated in Fig. 7.5. This benefits the quality of the training library. Next, rather than storing the data for each time step, the data is sub-sampled and only stored every two to four time steps while maintaining the same distribution of time steps (see Fig. 7.5). This allows more thermochemical conditions to be sampled for a given feasible training database memory footprint.

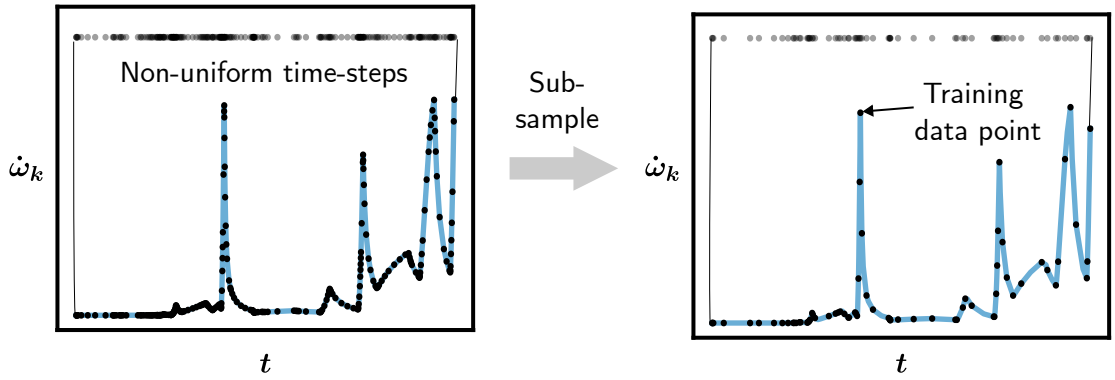


Figure 7.5: Schematic illustration of time step clustering and sub-sampling (plotted on a linear scale). Note, to help visualisation, only a fraction of the time steps are plotted

7.2.3 Thermochemical Phase Space

Visualising the phase space of the K-D thermochemical training database is not possible; therefore, Fig. 7.6 (left) instead shows a 3D projection of the full phase space with a single training simulation trajectory shown in red. It is useful to observe that there are correlations between species and that there is a clear manifold on which the training data points lie. This can be exploited later in the compression/reconstruction algorithm.

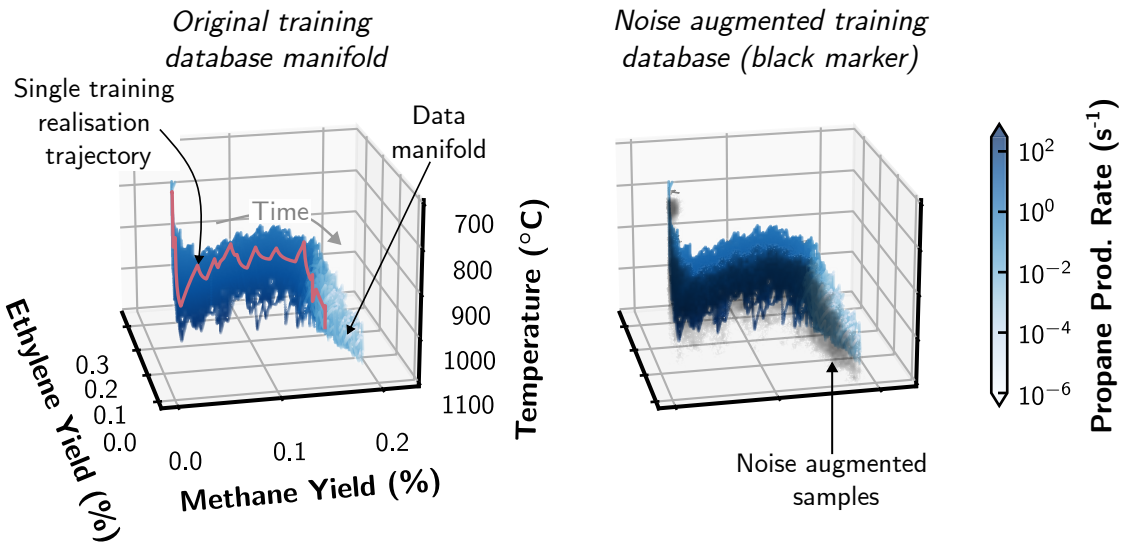


Figure 7.6: Scatter plot showing a projection of the high-dimensional thermochemical phase space for the training data, including noise augmented samples. Note that the number of training data points is downsampled by a factor of 20 and a log-scale used to help visualisation

Training Data Augmentation: Noise Injection

Since training simulations are performed for a 1-D homogeneous mixture, species diffusional effects are not taken into account. Diffusion acts similarly to applying synthetic perturbations to the thermochemical state [169] and pulling the phase-space trajectories away from the low-dimensional manifold shown earlier in Fig. 6.6. In order to artificially incorporate this effect into the training database, uncorrelated random noise is added to every existing data point Y_k, T & p already in the training database [169], as illustrated in Fig. 7.6. This noise perturbation, $\tilde{Y}_k = Y_k + \mathcal{N}(0, \sigma^2_{Y_k})$, is sampled from a multivariate Gaussian with variance (*e.g.*, $\sigma^2_{Y_k}$) equal to 5% of the maximum realistic range of the mass fraction for each species. These new samples (shown with black markers in Fig. 7.6) fill the space around the original manifold.

To sanitise the newly updated sample space, the constraint $\sum_{k=1}^{K} Y_k = 1$ is imposed and samples with an *exothermic* heat release rate $\dot{\omega}_T > 0$ are rejected, since this is unrealistic for most pyrolysis reaction regimes. The elemental/atomic carbon and hydrogen content is not constrained to be constant. This helps to provide support for interpolation even if the data points themselves are not explicitly accessed at inference time.

For the newly perturbed thermochemical state, the new values for $\dot{\omega}_k$, $\dot{\omega}_T$, etc. are evaluated from the reaction network. The final training database size is now double

the original unperturbed database. Since time integration is not required, generating noise-augmented samples only takes order of minutes.

7.2.4 The Training Database

The input state for the training database can be written as $\mathbf{X} \in \mathbb{R}^N \times \mathbb{R}^{K+2}$, where

$$\mathbf{X} = \begin{array}{c} \\ \text{time} \\ \downarrow \end{array} \overbrace{\longrightarrow}^{\text{state variables}} \begin{bmatrix} Y_1(t_1) & Y_2(t_1) & \cdots & Y_K(t_1) & T(t_1) & p(t_1) \\ Y_1(t_2) & Y_2(t_2) & \cdots & Y_K(t_2) & T(t_2) & p(t_2) \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ Y_1(t_N) & Y_2(t_N) & \cdots & Y_K(t_N) & T(t_N) & p(t_N) \end{bmatrix}, \tag{7.1}$$

and the output response $\mathbf{W} \in \mathbb{R}^N \times \mathbb{R}^{K+4}$ as

$$\mathbf{W} = \begin{bmatrix} \dot{\omega}_1(t_1) & \dot{\omega}_2(t_1) & \cdots & \dot{\omega}_K(t_1) & \dot{\omega}_T(t_1) & c_p(t_1) & c_v(t_1) & \bar{\mu}(t_1) \\ \dot{\omega}_1(t_2) & \dot{\omega}_2(t_2) & \cdots & \dot{\omega}_K(t_2) & \dot{\omega}_T(t_2) & c_p(t_2) & c_v(t_2) & \bar{\mu}(t_2) \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \dot{\omega}_1(t_N) & \dot{\omega}_2(t_N) & \cdots & \dot{\omega}_K(t_N) & \dot{\omega}_T(t_N) & c_p(t_N) & c_v(t_N) & \bar{\mu}(t_N)) \end{bmatrix}. \tag{7.2}$$

By applying the *frozen chemistry* assumption (see Sec. 6.2.3), other thermodynamic properties can be calculated from c_p and c_v. The temporal snapshots stored for each heat flux profile realisation are stacked vertically in Eqs. (7.1) & (7.2). Therefore, the training library maps all composition-dependent variables from the thermochemical input state to their output response $\{\mathbf{X} \mapsto \mathbf{W}\}$ (see Fig. 7.2).

7.3 Neural Network Architecture

Once we have compiled a training database, Fig. 7.7 show how machine learning is exploited to compress the high-dimensional thermochemical space onto a low-dimensional manifold using an encoder—$\phi : \mathbb{R}^K \mapsto \mathbb{R}^M$, where $M \ll K$—while simultaneously learning a regression mapping for the dynamic response in latent space $\dot{\boldsymbol{\omega}}_Z = \mathscr{F}_{NN,Z}(T, p, \boldsymbol{Z})$. Although not strictly correct, it is helpful to think of the latent space basis as a set of "meta"

or artificial species. During training, a mapping between the meta-species mass fractions $\boldsymbol{Z} = [Z_1, Z_2, \ldots, Z_M]^\top$, physical temperature & pressure, and the reaction rates projected into the latent space $\dot{\boldsymbol{\omega}}_Z = [\dot{\omega}_{Z_1}, \dot{\omega}_{Z_2}, \ldots, \dot{\omega}_{Z_M}]^\top$ is learnt (see Fig. 7.7). Through the set of source terms $\dot{\boldsymbol{\omega}}_Z$, this lumped representation of the dynamics is coupled with a convection-diffusion transport equation in a standard CFD flow solver. After time integration in reduced latent space has converged, the full multi-species spatio-temporal distribution of mass fractions in physical space $\widehat{Y}_k$ can be reconstructed during post-processing.

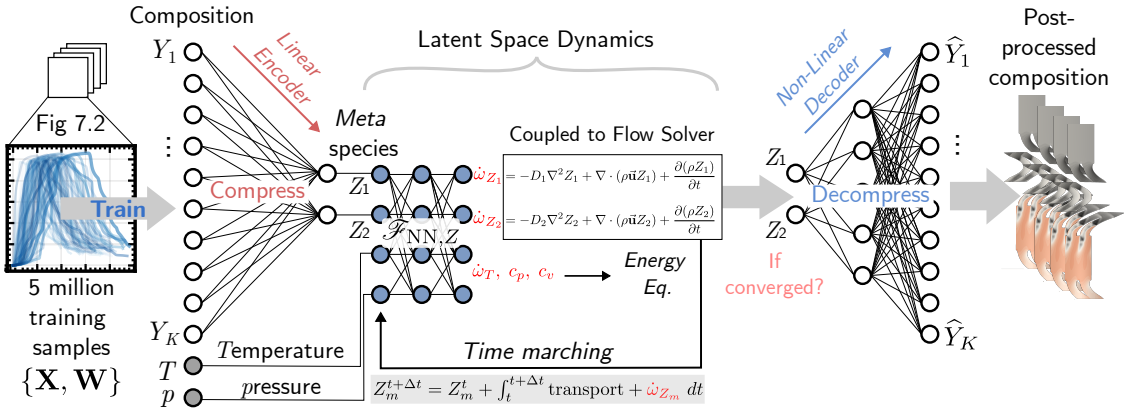


Figure 7.7: Schematic illustration showing the `ChemZIP` chemical autoencoder NN architecture for compressing the thermochemistry and coupling with a 3-D viscous CFD flow solver

The following subsections explain each of the individual components of the full `ChemZIP` workflow. Finer details about the neural network architecture and training procedure can be found in Appendix B.

7.3.1 Linear or Non-Linear Compression?

In Chapter 6 it was explained that the chemical reaction system is naturally attracted to a low-dimensional manifold (see Fig. 6.6). This means that there is redundancy in the underlying physics, which justifies the application of compression algorithms on the high-dimensional composition space. The compression mapping, or equivalently coordinate system transformation or projection, from composition space to latent space can either be linear or non-linear. To understand this transformation, we start with the species conservation equations, originally defined in Eq. (6.2), and rewritten in Eq. (7.3) in vector form. Equation (7.3) is then multiplied by the transformation/projection

operator $\mathbf{J}_\phi \in \mathbb{R}^{M\times K}$ on both sides:

$$\mathbf{J}_\phi \left(\frac{\partial \rho \boldsymbol{Y}}{\partial t} + \nabla \cdot (\rho \vec{\mathbf{u}} \boldsymbol{Y}) + \nabla \cdot (\boldsymbol{D}_{\text{eff.}} \nabla \boldsymbol{Y}) \right) = \rho \mathbf{J}_\phi \dot{\boldsymbol{\omega}} \tag{7.3}$$

where $\boldsymbol{D}_{\text{eff.}} = \rho \bar{\boldsymbol{D}} + \frac{\mu_t}{Sc_t} \in \mathbb{R}^{K\times K}$ is a square diagonal matrix of effective diffusion coefficients for all species, and

$$\mathbf{J}_\phi = \frac{\partial \boldsymbol{Z}}{\partial \boldsymbol{Y}} \equiv \left[\frac{\partial \boldsymbol{Z}}{\partial Y_1}, \ldots, \frac{\partial \boldsymbol{Z}}{\partial Y_K} \right] = f(\text{NN parameters, activation functions}), \tag{7.4}$$

represents the Jacobian of the linear or non-linear transformation mapping ϕ provided by the encoder of an AE or through the PCA matrix respectively, as shown in Figs. 7.7 & 7.8. Effectively, during training, a custom compression algorithm is learnt based on the given training data. In both cases, correlations between species are exploited to capture the majority of the data variance with only a few meta-species.

In the `ChemZIP` code, both linear and non-linear encodings have been implemented. The non-linear mapping is represented by an autoencoder-like neural network with multiple non-linearly activated hidden layers, while the linear mapping is a special case of an AE, where there are no hidden layers and the activation functions are linear. This is approximately equivalent to PCA. By formulating both as a neural network, it is easy to switch between them for different purposes. Moreover, because of the specialised algorithms now available to train NN, this formulation is significantly more scalable than standard PCA.

Linear *Encoding*

In the linear PCA [137] case, $\mathbf{J}_\phi$ is a $M \times K$ linear projection matrix composed of constant coefficients/weightings that are learnt during training. Effectively, the latent space meta-species become linear combinations of the high-dimensional physical species mass fractions, for example:

$$Z_1 = a_1 Y_1 + a_2 Y_2 + \ldots + a_K Y_K, \tag{7.5}$$

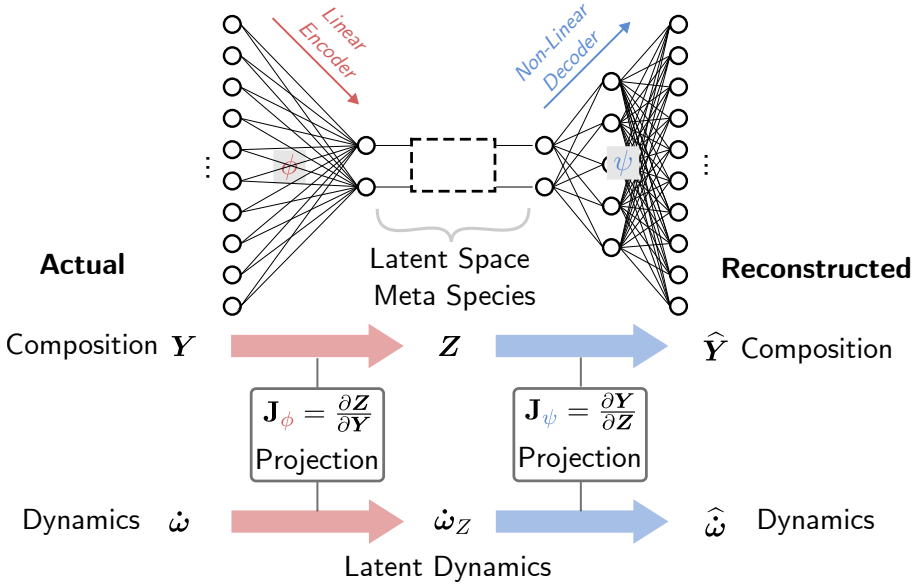


Figure 7.8: Schematic of isolated encoder-decoder architecture extracted from the full `ChemZIP` workflow in Fig. 7.7

where a_1 to a_K are the optimised NN weights (see Fig. 7.7). For the linear case, $\mathbf{J}_\phi$ is constant, does not vary in space & time, and is independent of the thermochemical state; therefore, Eq. (7.4) can be rewritten as a transport equation for the meta-species as follows,

$$\frac{\partial \rho \boldsymbol{Z}}{\partial t} + \nabla \cdot (\rho \vec{\mathbf{u}} \boldsymbol{Z}) + \nabla \cdot (\boldsymbol{D}_Z \nabla \boldsymbol{Z}) = \rho \dot{\boldsymbol{\omega}}_Z(T, p, \boldsymbol{Z}), \tag{7.6}$$

where $\dot{\boldsymbol{\omega}}_Z(T, p, \boldsymbol{Z}) = \mathbf{J}_\phi \dot{\boldsymbol{\omega}}$ (see Fig. 7.8), $\boldsymbol{Z} = \mathbf{J}_\phi \boldsymbol{Y}$ and $\boldsymbol{D}_Z \in \mathbb{R}^{M \times M}$ is the diffusivity matrix for the meta-species. This is equivalent to $\boldsymbol{D}_Z = \mathbf{J}_\phi \boldsymbol{D}_{\text{eff.}} \mathbf{J}_\phi^{-1}$, which simplifies to $\boldsymbol{D}_Z \equiv \left(\rho \bar{D} + \frac{\mu_t}{Sc_t}\right) \boldsymbol{I}$, where $\boldsymbol{I} \in \mathbb{R}^{M \times M}$ is the identity matrix, assuming that the molecular diffusion coefficient for all species is equal to $\bar{D}$ (as per Sec. 6.2.3). It is possible to write Eq. (7.4) in the form shown in Eq. (7.6) because $\mathbf{J}_\phi$ commutes with the convection and diffusion differential operators and therefore can pass under the $\nabla(\cdot)$, $\nabla \cdot (\cdot)$ and $\frac{\partial}{\partial t}(\cdot)$ operators.

Non-Linear *Encoding*

In the non-linear projection case, a fully-connected AE NN [137, 167] is used to extract low-dimensional structure in the physics. This provides greater levels of compression for a given accuracy requirement; however, in a multidimensional environment with convection and

diffusion, additional closure terms are required. This is because $\mathbf{J}_\phi \equiv \mathbf{J}_\phi(\boldsymbol{Y}) \neq$ constant and thus can vary in space and time as a function of composition. As a result, $\mathbf{J}_\phi$ does not commute with the convection and diffusion operators. Therefore, Eq. (7.6) is no longer valid and requires additional closure terms to balance it. Consequently, we restrict ourselves to linear projection.

7.3.2 Non-Linear Reconstruction

While the encoder is linear PCA, the decoder is a non-linear fully-connected AE-like NN. The physical species mass fractions can be decoded/reconstructed from low-dimensional space through the following transformation $\widehat{\boldsymbol{Y}} = \mathbf{J}_\psi \boldsymbol{Z}$ (see Fig. 7.8), where the Jacobian of the decoder $\psi : \mathbb{R}^M \mapsto \mathbb{R}^K$ is $\mathbf{J}_\psi = \frac{\partial \boldsymbol{Y}}{\partial \boldsymbol{Z}}$ and $\widehat{\square}$ signifies a reconstructed quantity (see Fig. 7.7). Despite using *linear encoding* for the reasons described above, applying *non-linear decoding* can condition the manifold while providing a smoother & higher-quality topology (see Fig. 6.9) [158], and can correct errors that would occur if the reconstruction were linear. In contrast with encoding, non-linear reconstruction is valid because it is only the variables that are decompressed, not the governing equations. Therefore, no closure terms are necessary. Although non-linearity increases the representational capacity, it is unclear what extra information is provided given that the manifold is linearly constructed (see Sec. 7.3.1). This should be further investigated.

7.3.3 Latent Space: Thermochemical Dynamic Response

In conjunction with learning a low-dimensional manifold, a non-linear regression mapping $[\dot{\boldsymbol{\omega}}_Z, \dot{\omega}_T, c_p, c_v, \bar{\mu}]^\top = \mathscr{F}_{\mathrm{NN},Z}(T, p, \boldsymbol{Z})$ (see Fig. 7.7) is learnt for the dynamic response of the reaction system projected into latent space (see Fig. 7.8 and Eq. (7.6)) for a given the state $T, p, \boldsymbol{Z}$. A fully-connected feed-forward neural network—a highly rich and flexible function approximator—is chosen over other methods used in the literature, such as Gaussian processes, which are used by Malik *et al.* [157]. Neural networks are preferred because of their scalability to a much larger corpus of training data than the other methods. Since Gaussian processes and other such non-parametric kernel methods

effectively perform inference through a weighted average over *all* points in the training dataset, training and inference times scale poorly $\mathcal{O}(N^2)$ to large datasets.

It was decided to learn the source terms $\dot{\boldsymbol{\omega}}_Z$ (from Eq. (7.3)) rather than the solution propagator $\boldsymbol{Z}(t) \rightarrow \boldsymbol{Z}(t + \Delta t)$, as discussed in Sec. 6.3.2 and illustrated in Fig. 6.8, as it is simpler to couple with the CFD solver (since operator splitting is not required) and does not place restrictions on the time step size. Furthermore, notice that the NN for the reaction heat absorption rate and the thermophysical properties are learnt as functions of the meta-species as opposed to the physical species mass fractions.

During the development of the `ChemZIP` workflow, three key design decisions were made to improve the performance of `ChemZIP`. First, it was decided to incorporate the temperature and pressure dependency of directly as inputs to $\mathscr{F}_{\text{NN},Z}(T, p, \boldsymbol{Z})$ mapping (see Fig. 7.7), rather than passing T & p through the encoder layer. This simplifies and accelerates coupling with the flow solver because the thermochemical state does not need to be repeatedly encoded and decoded at every time integration step to follow changes in temperature. Encoding is done once at the start of a simulation and decoding is done once at the end of a simulation.

Second, rather than independently training the NN parameters for the low-dimensional manifold mapping from the parameters for the latent-space dynamic response mapping, they are both trained concurrently. This coupled training procedure helps to avoid complications with a non-unique and non-smooth manifold topology discussed in Ref. [170] and illustrated in Fig. 6.9. The combined loss function naturally penalises poor manifold quality as the fitting error for $\mathscr{F}_{\text{NN},Z}(T, p, \boldsymbol{Z})$ implicitly increases.

Finally, a purely data-driven approach is adopted to approximate $\mathscr{F}_{\text{NN},Z}(T, p, \boldsymbol{Z})$. While other researchers have chosen to physically constrain the latent space reaction dynamics in the form of the Arrhenius rate law and the law of mass action [162], we have opted for a NN approximator. This is because early investigations indicated that the representational capacity and the number of trainable parameters of the physics-inspired approach were insufficient, resulting in an intolerably low level of accuracy. For example, with three meta-species, a physics-inspired latent dynamics approach would only have around 21 parameters to fit, which is not the case for a purely data-driven approach. It is

felt that although this is less interpretable, there is no physical basis for the latent space dynamic response to be of the same functional form as the physical reaction space.

Recalling that low-latency inference is at the heart of `ChemZIP`, the number of neurons & the number of hidden layers are kept to a minimum and separate NN approximators are used for the following outputs $\dot{\boldsymbol{\omega}}_Z, \dot{\omega}_T, c_p, c_v, \bar{\mu}$. Therefore, smaller NNs can be used, which reduces the inference time. The detailed neural network architecture is defined in Appendix B.

7.3.4 The Loss Function

The trainable NN parameters (*i.e.,* weights and biases) of the encoder, decoder and non-linear regression for the system dynamics are optimised using enhanced version of stochastic gradient descent [137]. The NN model was built and subsequently trained using the KERAS [171] and TENSORFLOW v2.14.0 [172] APIs. The first step is to define a loss function to be minimised (*i.e.,* $\arg\min_{\boldsymbol{\theta}} \mathscr{L}$). This is defined as follows $\mathscr{L} = w_1\mathscr{L}_{\text{chemistry}} + w_2\mathscr{L}_{\text{thermo}}$, where

$$\begin{aligned}\mathscr{L}_{\text{chemistry}} = \underbrace{\lambda_1 loss(\widehat{\boldsymbol{Y}} - \boldsymbol{Y})}_{\text{reconstruction}} + \underbrace{\lambda_2 loss\big(trans(\dot{\boldsymbol{\omega}}_Z) - trans(\mathbf{J}_\phi \dot{\boldsymbol{\omega}})\big)}_{\phi\text{: encoded dynamics}} \\ + \underbrace{\lambda_3 loss\big(trans(\dot{\boldsymbol{\omega}}) - trans(\mathbf{J}_\psi \dot{\boldsymbol{\omega}}_Z)\big)}_{\psi\text{: decoded dynamics}}\end{aligned} \tag{7.7}$$

in which $loss(\boldsymbol{Y})$ is defined as $loss(\boldsymbol{Y}) = \sum_{n=1}^{N}\sum_{k=1}^{K} \log\left(\cosh\left(Y_{k,n}\right)\right)/(NK)$ and

$$trans(x) = \text{sign}(x + 0.0001)\sqrt{|x + 0.0001|}$$

[158]. The second component of $\mathscr{L}$ is defined as

$$\mathscr{L}_{\text{thermo}} = \text{MSE}(\widehat{c}_p - c_p) + \text{MSE}(\widehat{c}_v - c_v) + loss(trans(\widehat{\dot{\omega}}_T) - trans(\dot{\omega}_T)), \tag{7.8}$$

where $\text{MSE}(\mathbf{X}) = \sum_{n=1}^{N}(x_n - \hat{x}_n)^2/N$ is the mean square error. The structure of the loss function can be broken down as follows.

- *Reconstruction:* the reconstruction loss is the information lost from encoding $\rightarrow$ decoding.
- *Encoded dynamics*: an error signal is generated between the actual net species production rates projected into latent space and those predicted by $\mathscr{F}_{\text{NN},Z}(T, p, \boldsymbol{Z})$.
- *Decoded dynamics*: an error signal is generated between the predicted latent space dynamics projected into physical space and the actual net species production rates.
- $trans(\cdot)$: training NNs is challenging when the outputs vary over many orders of magnitude. Since this is the case for $\dot{\boldsymbol{\omega}}$ & $\dot{\omega}_T$ (see Fig. 7.4), they are power-transformed (see Fig. 7.9(a)) to prevent training from skewing toward extreme values by stabilising the statistical variance [146], as shown in Fig. 7.10. Errors at low orders are as important as those at high orders. This transformation tends to linearise the input-output relationship, making it easier to learn.
- $loss(\cdot)$: the log-cosh loss function [173] is used for similar reasons as above. A MSE loss function would bias the model strongly towards high magnitude values for $\dot{\boldsymbol{\omega}}$ & $\dot{\omega}_T$, whereas log-cosh provides a more balanced treatment across a range of orders of magnitude (see Fig. 7.9(b)).

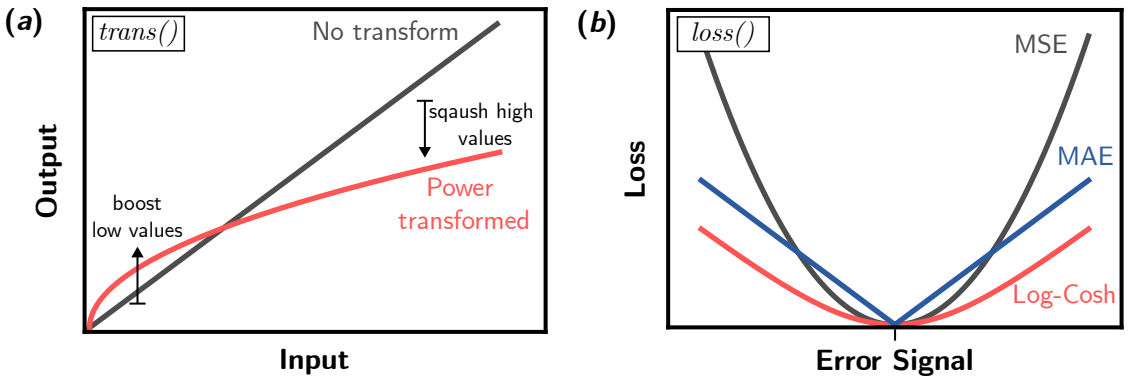


Figure 7.9: **(*a*)** Schematic illustration of the power transform pre-processing function and **(*b*)** schematic illustration of various loss functions

The multi-scale nature of some of the dependent variables is a critical challenge for the learning algorithm [146]. For more details on training the ML model and for values used for the hyperparameters, the reader is referred to Appendix B. In the current example, training takes approximately 12 to 24 hours for convergence.

An alternative pathway to defining and optimising the loss function would be to optimise over time-integrated multi-step model predictions. However, it is chosen to avoid this approach because error cancellation over the integration could lead to misleadingly low values for the loss function, even if single-step errors are large. This would not lead to a robust model when deployed in an unfamiliar 3-D viscous flow environment.

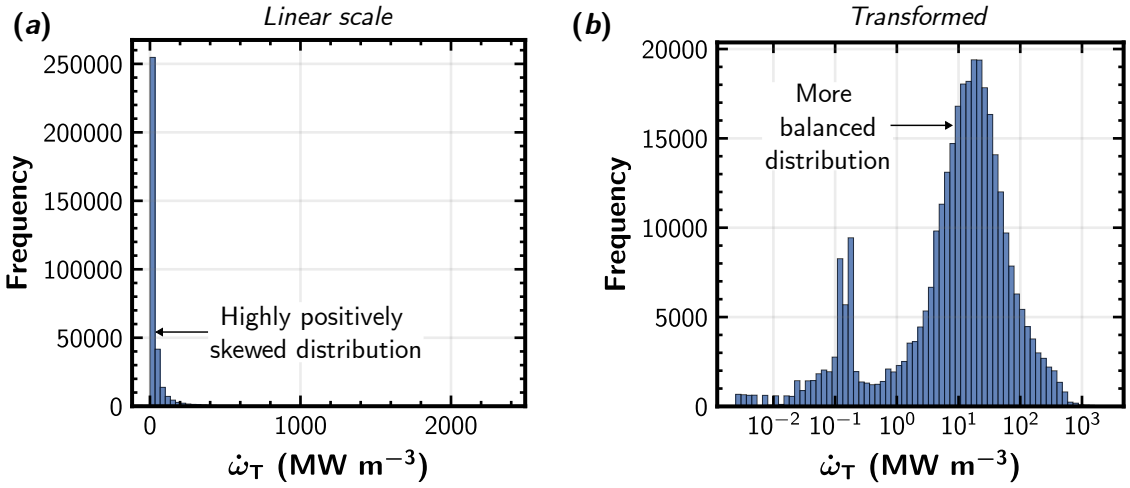


Figure 7.10: A histogram of the training database for the reaction heat absorption rate plotted on **(*a*)** a linear scale and **(*b*)** a logarithmic scale

Physical Constraints

An option has been added to encourage the enforcement of physical constraints [151] including elemental mass conservation. This constitutes adding additional penalising terms to the loss function. Although incorporating prior knowledge can improve generalisation, we have found that it can restrict the flexibility of the network and reduce accuracy.

Adjusting for Training Data Skew: Sample Weightings

For each sample in the training database $\mathbf{X}$, its contribution to the loss function $\mathscr{L}$ is weighted by a sample weight during training. When all sample weights were set to

1.0, it was found that accuracy was skewed towards high-heat-flux conditions because many more time steps were taken by CANTERA compared to low-heat-flux conditions in which the gradients are low. Therefore, the distribution of training data was skewed to regions of high dynamic response. To compensate for this, samples generated under low-heat-flux conditions are weighted higher, inversely proportional to the number of time steps taken for a given heat flux profile realisation.

7.3.5 Data Pre-Processing Pipeline

There are two main steps to pre-processing before training, culling unimportant species and scaling the inputs and outputs.

Species Cull

The first step in the pre-processing pipeline is to eliminate minor species with low variance across the training dataset using `VarianceThreshold()` from SCIKIT-LEARN. These are typically fast-reacting intermediate radical species whose concentration distributions we are not interested in predicting. This makes training easier. Typically, this process reduces the number of species by a factor of 3, but this depends on the mechanism. As a result, training speeds up, and since the highly non-linear radical species are removed, manifold topology also improves and practically eliminates numerical stiffness as we will see in Sec. 8.3.2. Furthermore, since the learning algorithm does not have to predict radical mass fractions, which are on the order of $<10^{-6}$, the accuracy of higher mass fraction species improves.

Data Normalisation

It is standard practice to normalise the different inputs (*e.g.,* temperature and mass fractions) and outputs on the same scale of [0, 1] in order to improve the performance of loss function optimisation. In many approaches in the literature, the mass fraction of each species is individually scaled, so that every species ranges from [0, 1]. In contrast, in this work, the entire mass fraction vector is shifted and scaled by the *same* factor, as follows $\boldsymbol{Y}_{\text{scaled}} = \frac{\boldsymbol{Y} - Y_{\text{min.}}}{Y_{\text{max.}} - Y_{\text{min.}}}$. This slightly biases the training to species with a higher mass fraction, which is desired because these are typically the most important species.

7.3.6 Error Metrics

To assess the accuracy and generalisability of the trained `ChemZIP` model on N_{test} unseen test data samples, the R^2 score, the normalised median absolute error (NMdAE) and the normalised root mean square error (NRMSE) metrics are used to measure goodness of fit. The R^2 score measures the amount of variance in the data that is a consequence of the true variability in the quantity of interest. Taking the reconstruction of the k^{th} species mass fraction as an example, these three metrics are defined as

$$R^2 = \left(1 - \frac{\sum_n^{N_{\text{test}}} (Y_{k,n} - \widehat{Y}_{k,n})^2}{\sum_n^{N_{\text{test}}} (Y_{k,n} - \overline{Y}_k)^2}\right), \tag{7.9}$$

$$\text{NMdAE} = \frac{\text{Median}\left(\left|Y_{k,1} - \widehat{Y}_{k,1}\right|, \ldots, \left|Y_{k,N_{\text{test}}} - \widehat{Y}_{k,N_{\text{test}}}\right|\right)}{\left|\overline{Y}_k\right|}, \tag{7.10}$$

and

$$\text{NRMSE} = \frac{\sqrt{\sum_n^{N_{\text{test}}} (Y_{k,n} - \widehat{Y}_{k,n})^2}}{\left|\overline{Y}_k\right| \cdot N_{\text{test}}} \tag{7.11}$$

where $Y_{k,n}$ is the ground truth, $\widehat{Y}_{k,n}$ is the prediction for the k^{th} species & n^{th} sample, and $\overline{Y}_k = \frac{1}{N_{\text{test}}} \sum_n^{N_{\text{test}}} Y_{k,n}$. Note that the R^2 metric ranges between 0% and 100%, where 100% represents a perfect prediction.

No regression metric is perfect for fairly condensing inaccuracies across the full error spectrum into a single scalar value. For example, for variables that vary over many orders of magnitude, the NRMSE disproportionately penalises errors at high magnitudes. Therefore, we tend to prefer the NMdAE, but this is also not ideal. Therefore, it is better to report several metrics.

7.4 Latent Phase Space Manifold

Figure 7.11 shows a visualisation of the high-dimensional training data projected into low-dimensional latent space—$\boldsymbol{Z} = \mathbf{J}_\phi \boldsymbol{Y}$ (see Fig. 7.8). The latent phase space is defined by Z_m and the temperature T. In this example, the bottleneck layer is limited to 2 latent meta-species Z_1 & Z_2 to enable visualisation. One simplified way of interpreting the latent space is that it is a filtered version of the physical mass fraction vector. It is formed by

eliminating redundant and correlated species. It can be seen that the manifold (consisting of 4 million training data points) is non-overlapping and that the reaction heat absorption rate varies smoothly across it without suffering from issues of non-uniqueness. This indicates robustness, underlining the benefits of simultaneously learning the encoder and the non-linear dynamics regression. Interestingly, for this particular example, Fig. 7.11 suggests that $\dot{\omega}_T$ is more strongly correlated with temperature and that the meta-species have a weaker effect. The manifold in Fig. 7.11(a) thin, which can pose challenges when deploying the ChemZIP model and may need to be addressed in future work.

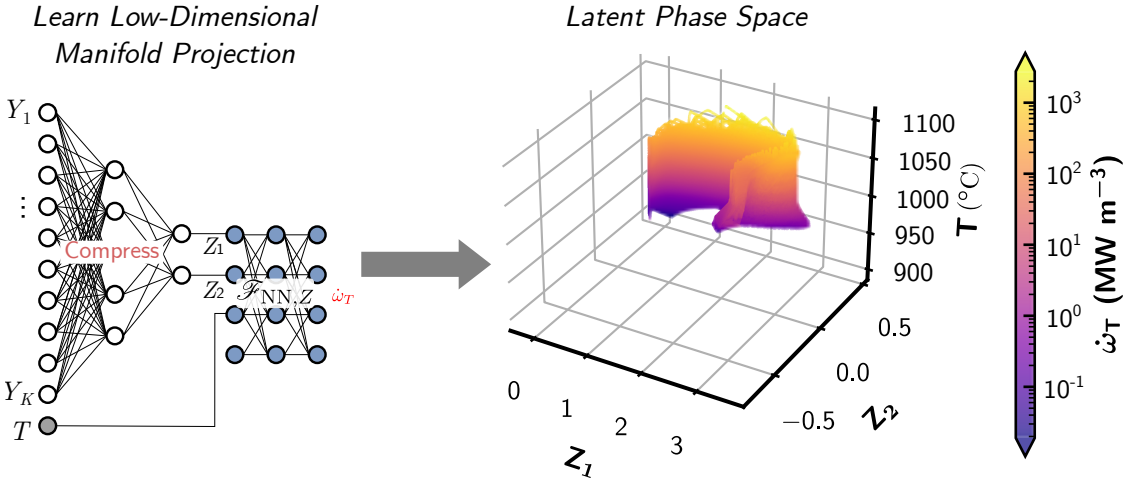


Figure 7.11: Latent space manifold for the training database constructed from Z_1, Z_2 & T and coloured by the reaction heat absorption rate $\dot{\omega}_T$ as the dependent variable (on a log-scale)

7.5 Coupling ChemZIP with a Flow Solver

The trained ChemZIP model written using the TensorFlow (TF) API can be deployed into any CFD flow solver relatively easily by executing the steps summarised below and illustrated in Fig. 7.12.

1. For each of the M meta-species (where $M \sim 2-4$), add a convection-diffusion transport equation of the form shown in Eq. (7.6) using standard discretisation methods already built into the solver.

2. To update the source terms $\dot{\boldsymbol{\omega}}_Z$ at each time iteration given new values for $T, p, \boldsymbol{Z}$, the TF model for $\mathscr{F}_{\mathrm{NN},Z}(T, p, \boldsymbol{Z})$ must be queried. This happens in each mesh cell.

3. Add meta-species wall boundary condition, which is a zero diffusive flux condition projected onto the latent space, *i.e.*, $\mathbf{J}_\phi \nabla \boldsymbol{Y}|_{\text{wall}} \equiv \nabla \boldsymbol{Z}|_{\text{wall}} = 0$ [174].

4. Add additional heat absorption sink term $\dot{\omega}_T$ to the energy equation, as shown in Eq. (6.4) in Chapter 6. This source term is the dominant mechanism driving interactions between the aerothermodynamics and the chemistry.

5. Clip the inputs to the NNs within known ranges set by the training data space to prevent dangerous model extrapolation. For example, currently, the temperature input into the model is clipped to the range 600 °C to 1150 °C.

6. Fluid properties c_p & c_v are now set as a function of T & $\boldsymbol{Z}$.

Effectively, all complex reaction effects are lumped into a set of source terms and changes in fluid properties. Crucially, as will be demonstrated in Chapter 8, the stiffness (see Fig. 6.2) of the reaction system is alleviated, which means that the latent chemistry does not impose any time-step restrictions on the datum flow solver. Therefore, the basic *explicit* time integration scheme, multigrid acceleration, etc., can be used without modification, significantly accelerating convergence.

This methodology is agnostic to the level of fidelity of the flow solver, which could be anything from RANS/URANS to LES. Currently, `ChemZIP` is coupled with the industry-standard turbomachinery flow solver Tblock [175]. However, it could also be integrated with a commercial solver such as Fluent [176] by adding a module using user-defined functions, or combined with a 2D throughflow solver for fast iterative blade and multi-stage design.

Currently, inference is done only on the CPU, as it is the only hardware for which Tblock is written. For parallel simulations, a separate TF model is instantiated for each `MPI` process. However, since inference is essentially a series of matrix-vector multiplications, it is much faster on GPUs than on CPUs. Therefore, a flow solver capable of running on GPUs, as well as handling complex geometries, was developed as part of this work. This flow solver, called U-Tblock, is based on the same numerical methods developed for the structured-mesh code Tblock but re-engineered to work for an unstructured mesh. The necessary code paths for both CPU- and GPU-based parallelisation are automatically

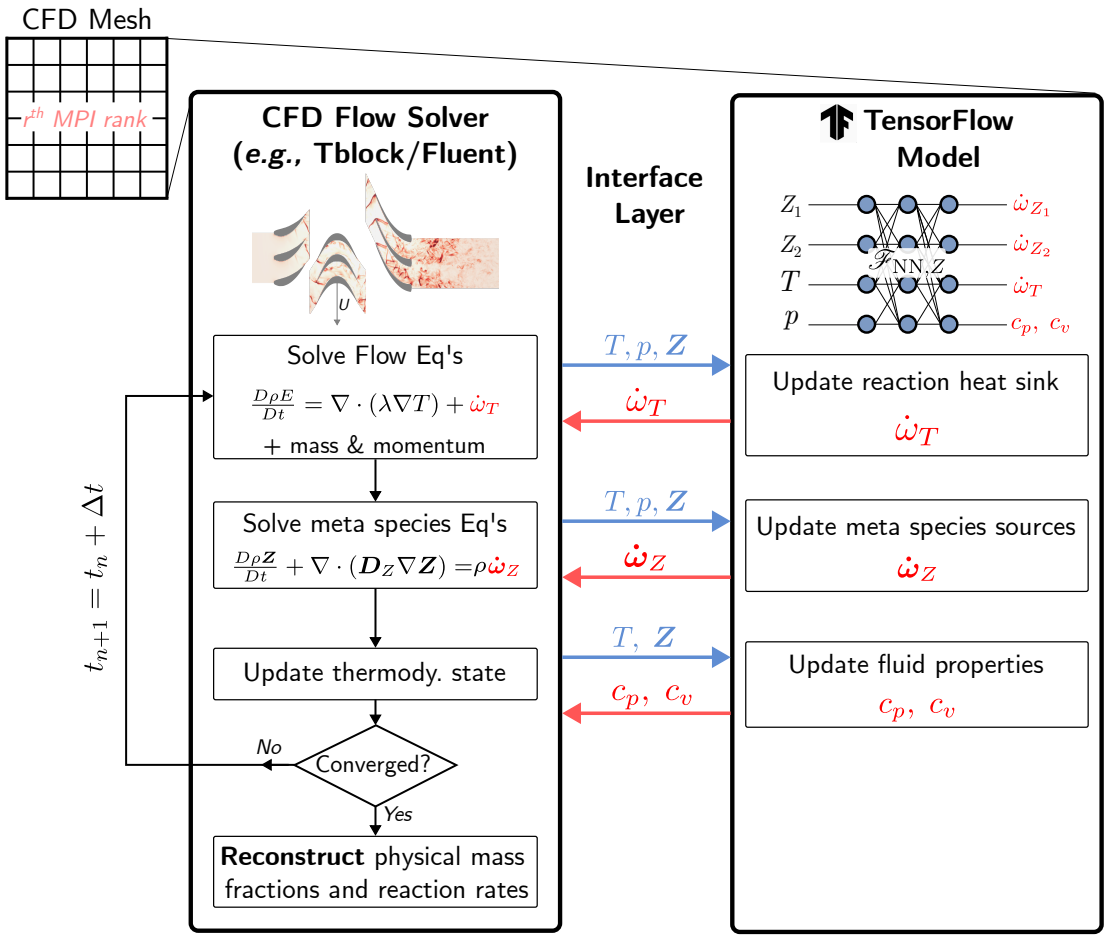


Figure 7.12: Flow chart of the coupling process between the CFD flow solver and the trained `ChemZIP` model written using the TensorFlow API

created via a templated code generation approach using the OP2 domain-specific language (compiler) [177]. Basic verification of the solver is illustrated in Fig. 7.13. This code is still in the verification stage and so will not be presented in detail in this thesis. In future work, `ChemZIP` will be integrated into U-Tblock.

7.5.1 Tblock–`ChemZIP` Interface

One challenge is that the in-house CFD solver Tblock is written in `Fortran` while the TF NN model is most easily accessible through the high-level `Python` interface. Therefore, an interface layer is used to efficiently communicate large multidimensional arrays between Tblock (`Fortran`) and TensorFlow (`Python`). For each required field (*e.g.*, temperature), the entire 3D array equal to the mesh size is shared at each time step per `MPI` rank. This is much more efficient than looping over the mesh and sharing

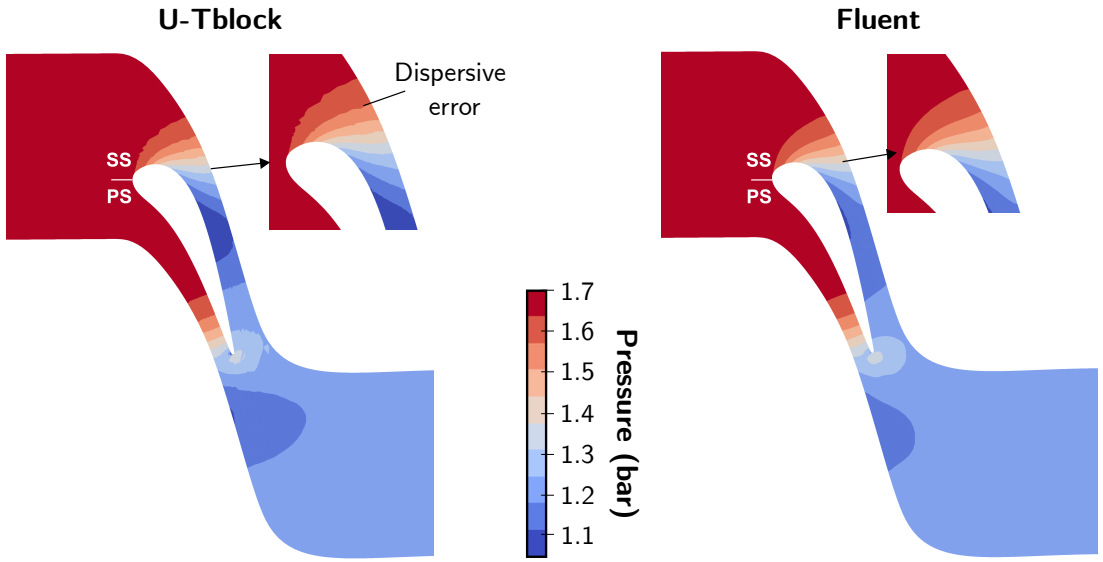


Figure 7.13: A verification of pressure distribution prediction on a highly-loaded gas turbine nozzle guide vane from U-TBLOCK against FLUENT

data at each loop iteration. This is because code coupling and TF API overheads are amortised, and evaluating NNs is much more efficient on larger batches of data.

The first step in the interface layer is to pass a reference to the data stored in memory between `Fortran` ↔ `C`. This allows the use of the *C Foreign Function Interface* (CFFI) API to efficiently share data between `C` ↔ `Python`. In this process, only data references are shared across the interface, which avoids the overhead of `memcpy()` operations. Where the data layout needs to be reshaped before/after data-sharing, persistent buffers are created once at the start of the program to avoid dynamic memory allocation overheads. Although there is a small additional overhead interfacing `Fortran` ↔ `C` ↔ `Python`, the flexibility of the `Python` interface helps accelerate the development and testing process.

Before using the TF model, it is converted to a TFLITE model. This is a lightweight version of the original TF API that is optimised for low-latency inference. The backend of TFLITE is written in `C++`; therefore, in the future, the TF model can be more directly interfaced with TBLOCK (*i.e.,* `Fortran` ↔ `C++`), as demonstrated in Ref. [178].

7.6 Summary

In this chapter, the new `ChemZIP` methodology has been introduced to computationally accelerate 3D viscous chemically reacting flow simulations by learning from a comprehen-

sive and intelligently curated training database. A methodology has been introduced to mitigate the training-data-generation bottleneck that frequently impedes the efficacy of ML approaches. Rather than taking weeks, a large suite of training data can be generated in only a few hours. We have explained how to (1) compress the high-dimensional composition space vector onto a low-dimensional manifold using machine learning (2) learn a fast mapping between the new latent space meta-species and their corresponding dynamics and (3) efficiently and simply couple a set of meta-species transport equations with a CFD flow solver of arbitrary fidelity. Although it has not yet been verified whether the data-driven `ChemZIP` approach can outperform state-of-the-art acceleration approaches for combustion *i.e.,* FGM, this approach is certainly equally as fast and, most importantly, significantly more universal. Generality beyond combustion (*e.g.,* for reactor kinetics modelling) is the primary motivation for developing this methodology. In many reacting flow scenarios, such as hydrocarbon cracking reactions, there is no expert knowledge to define the latent space state vector (which would be the mixture fraction and reaction progress variables in the case of FGM). A data-driven solution has helped to overcome this, learning the reduced basis automatically from the training data. Moreover, it is highly likely that `ChemZIP` could improve on FGM by learning a more appropriate basis vector from the data by exploiting non-linearity. In other data-driven approaches for chemistry, non-linear reconstruction is seldom harnessed for 3D reacting flow test cases. The next section will verify the performance and accuracy of `ChemZIP` in a hierarchy of different problems.

Chapter 8

Verification, Benchmarking and Proof of Concept

8.1 Introduction

To assess the accuracy, robustness and generalisability of `ChemZIP`, it is compared against the FOM for both 1D PFR and 3D viscous reacting flow simulations. The *test dataset*, which comprises unseen data withheld from the training dataset, is used to perform this evaluation. Generally, assessing the accuracy and reliability of a computational model can be achieved by either performing *verification* or *validation*. For verification, the so-called "ground truth" is a known numerical solution, which in this case is a reaction simulation with a complete detailed kinetic model (*i.e.,* the FOM). For validation, the ground truth is provided by experimental data. In this work, a verification assessment is more appropriate since it helps to isolate errors caused by the modelling of the reaction network rather than those due to other effects. This prevents us from misinterpreting errors that are not related to `ChemZIP`; for example, those due to uncertainties in the kinetic model (which can always be replaced with a more accurate one).

Verification is carried out in three phases.

1. *Input-output "a priori" testing*: the reconstruction error for the mass fractions, production rates and reaction heat absorption rate is assessed purely based on input-output pairs without time integration and thus are not subject to error accumulation.

2. *Time-integrated "a posteriori" 1D PFR testing*: `ChemZIP` is incorporated into the 1D PFR solver and is time-integrated. The accuracy over the full time history is assessed. This is important to determine the robustness of the model over thousands of time steps, where errors can accumulate. While the heat flux profiles imposed have not previously been directly seen, they come from the same statistical distribution that generated the training data and are thus still in the interpolative regime.

3. *3D viscous reacting flow CFD*: the model is then deployed in a 3D viscous flow solver and compared against a state-of-the-art reacting flow solver, *i.e.,* FLUENT. Since the flow environment is entirely different from the training data set, this is an "out-of-distribution" test case and is a good assessment of reliability and robustness. Although many researchers have demonstrated successful data-driven methodologies for accelerating chemical kinetics in a 1-D environment, far fewer have actually deployed and verified their models into a fully-fledged CFD code.

4. *Proof of concept*: finally, as a proof of concept, we demonstrate that it is now possible to predict the spatial distribution of yields, reaction rates and coking rates along the bladed path of the turbo-reactor.

Throughout this chapter, the kinetic model used is a 65-species propane-steam pyrolysis mechanism typical of industrial steam cracking. For now, the mechanism is small to accelerate the development time; however, an arbitrarily complex mechanism can be used as a drop-in replacement in the future "application phase" of this work. Although this is a relatively simple mechanism, it is significantly larger than the ones used by other practitioners developing surrogate models to accelerate chemical kinetic simulations (*e.g.,* in Refs. [146, 161]).

8.2 Verification: Error Statistics

The following sections present a hierarchy of verification test cases in which `ChemZIP` is compared against ground-truth test data. In all cases, the ground truth is the FOM with a complete detailed kinetic model (in this case, consisting of 65 species). The error metrics reported here are defined in Eqs. (7.9)–(7.11) in Sec. 7.3.6. For dependent variables that can rapidly vary over many orders of magnitude—for example, the reaction heat absorption rate $\dot{\omega}_T$—a logarithmic scale (defined only for strictly positive numbers) or arcsinh scale (which is well defined for all real numbers) is used to improve visualisation. It can be assumed that predictions for quantities of the physical space are in fact reconstructions from the latent space (denoted by $\widehat{\square}$) calculated during post-processing.

For all test cases below, it is assumed that the FOM model follows the governing equations and modelling assumptions laid out in Sec. 6.2. These assumptions also apply to `ChemZIP`. For example, the cross-species heat flux term is neglected in both. These assumptions help to isolate errors in the model, but can be refined in future work.

It should be noted that the verification of `ChemZIP` shown in this chapter is a proof of concept. Detailed tuning of the hyperparameters & NN architecture, and intelligent enhancement of the training database are left for future work.

8.2.1 A Priori Input-Output Reconstruction

For the test dataset, we start by assessing the input-output relationship without time integration. The test dataset consists of 500,000 data points corresponding to approximately 1000 different heat flux profile realisations withheld from the training dataset. This data is generated in the same way as described in Sec. 7.2. The results are presented in the form of parity plots, where the red dashed 45° line (see Fig. 8.1) represents perfect reconstruction. To evaluate the test accuracy for the source terms $\widehat{\omega}_k$ & $\widehat{\omega}_T$, the thermochemical state is

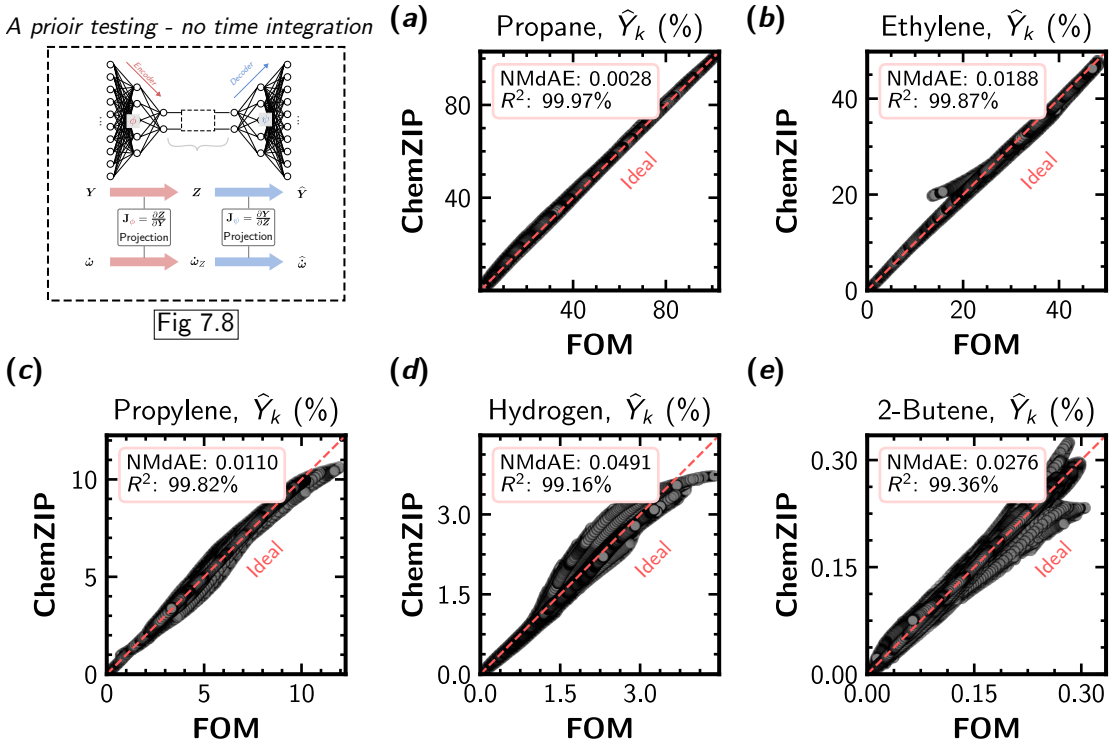


Figure 8.1: A priori (*i.e.*, no time integration) parity plot (500,000 data points) comparing reconstructed dry yields predicted by `ChemZIP` against the ground truth FOM data for 5 species **(*a*)**–**(*e*)**

first projected into latent space, the non-linear regression map is then queried, and finally the predicted values for the dynamics are projected back into the physical space.

Figure 8.1 shows that for both major (*i.e.,* high concentration) and minor (*i.e.,* low concentration) species, the loss of information in the compression-decompression pipeline is low for the species yields. The errors are highest for minor species like hydrogen (see Fig. 8.1(d)), but the R^2 score is still above **99%**. Figure 8.2 suggests that the regression fits for the source terms and thermodynamic properties are fairly good over a wide range of conditions, with the R^2 score above **97%**. Again, hydrogen exhibits the worst

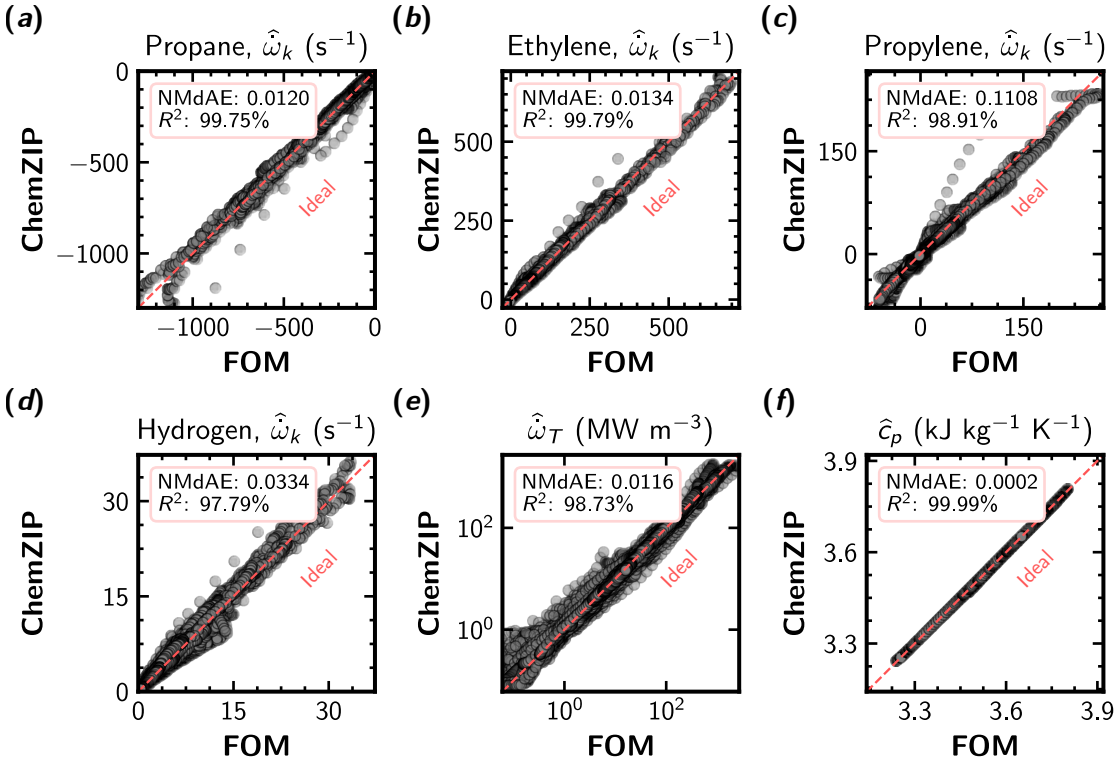


Figure 8.2: A priori (*i.e., without* time integration) parity plot (500,000 data points) comparing reconstructed `ChemZIP` predictions against the ground truth FOM data for **(*a*)**–**(*d*)** the net production rates **(*e*)** the heat absorption rate (on a log scale) and **(*f*)** the heat capacity

performance. This is likely because its magnitude is several decades lower than that of the other species, indicating some refinement of the loss function (see Sec. 7.3.4) and additional data pre-processing is required to counteract this bias. However, generally this hierarchy of accuracy between major and minor species is tolerable as major species are more valuable and more important to accurately predict. In several regions there is a significant deviation from the ideal 45° (see Figs. 8.1 & 8.2) line. This is a consequence of insufficient

training data in some regions of the thermochemical phase space. For the reaction heat absorption rate (see Fig. 8.2(e)), values with a lower magnitude appear to have the most scatter. This is partly due to lack of training data coverage (which can be inferred from Fig. 7.10(b)) and partly due to challenges with optimising the NN parameters for multi-scale data. However, this should not have a significant impact on the overall predictions of the system since the rates of reaction are relatively insignificant in this regime.

Table 8.1 summaries the error metrics for the training and testing datasets. The values do not differ substantially between the training and testing datasets, indicating that the model is generally not significantly overfitting. However, the training and testing datasets were produced from the same statistical distribution, so this is not a considerably challenging assessment. There appears to be some overfitting for the reaction heat absorption rate, highlighted by a notable increase in NRSME between train and test set, as well as for the net production rates, indicated by the drop in the R^2 score. The latter is a consequence of poor predictions for the less important minor species (*i.e.,* $Y_{k,\max.} < 5\%$), which contaminates the average metric. Fortunately, for the more important major species, the predictions of net production rates are much better (see Fig. 8.2).

Table 8.1: Error metrics for the training dataset and the test dataset for both single-step (a priori) and multi-step (a posteriori) predictions

Variable	$\frac{1}{K}\sum_k^K \text{Metric}(\widehat{Y}_k)$	$\frac{1}{K}\sum_k^K \text{Metric}(\widehat{\dot{\omega}}_k)$	$\widehat{\dot{\omega}}_T$
Training: R^2 (%)	99.1590	99.0813	99.2982
A Priori Test: R^2 (%)	99.5187	98.2051	98.6606
A Posteriori Test: R^2 (%)	98.9618	95.1290	98.7629
Training: NMdAE (−)	0.01348	0.01437	0.01135
A Priori Test: NMdAE (−)	0.01496	0.01568	0.01099
A Posteriori Test: NMdAE (−)	0.01488	0.01340	0.00931
Training: NRMSE (−)	0.04154	0.22035	0.17354
A Priori Test: NRMSE (−)	0.03059	0.23628	0.23044
A Posteriori Test: NRMSE (−)	0.03886	0.22654	0.22753

8.2.2 A Posteriori Multi-Step 1-D PFR Predictions

Although input-output accuracy is a necessary condition for good performance when deployed in CFD solver, it is not sufficient. It does not address the sensitivity of the model

to error accumulation in a realistic time-integrated online environment. Therefore, to now assess the reliability of `ChemZIP` deployed in a more realistic setting, it is time-integrated in a 1D PFR environment under a complex set of imposed heat flux distributions (described in Sec. 4.1.1). The trained model is integrated using an explicit $4^{\text{th}}/5^{\text{th}}$-order adaptive Runge-Kutta method embedded within `solve_ivp()` from SciPy [179]. The latent space reaction system, consisting of only three meta-species transport equations defined in Eq. (7.6), is integrated using a coupling procedure similar to that shown in Fig. 7.12. The very fact that an *explicit* method works (rather than a stiff *implicit* method), implies that the numerical stiffness of the reaction system has been alleviated, dramatically accelerating convergence. This will be elaborated on later.

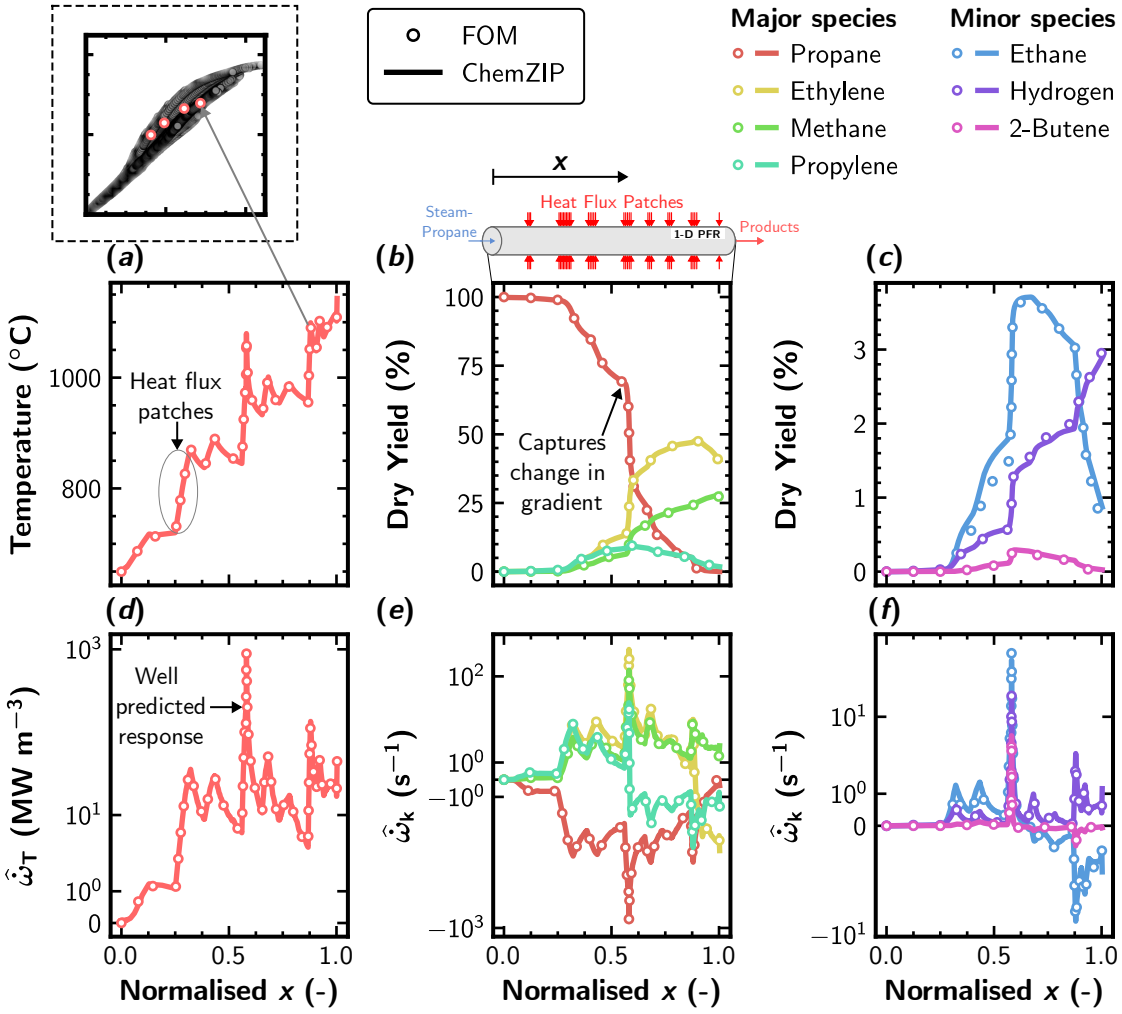


Figure 8.3: 1D PFR time integrated solution history for **(*a*)** temperature **(*b*)**–**(*c*)** major/minor species dry yields **(*d*)** the reaction heat absorption rate and **(*e*)**–**(*f*)** the major/minor species net production rates. Note that the circular markers denoting the FOM are significantly downsampled for clarity

Figure 8.3 shows the solution history for a single heat flux profile realisation while Fig. 8.4 condenses the time histories for all 1000 1D PFR simulations from the test dataset into a series of parity plots. Assessing the performance of `ChemZIP` when subjected sharp gradients in temperature, as seen in Fig. 8.3, ensures representative conditions for the turbo-reactor are covered. It is clear that even for a complex external heat load (see Fig. 8.3(a)), accurate production rates can be reconstructed from latent space (see Fig. 8.3(d))–(f)). As a result, following a heat flux perturbation, the temporal gradient in the yield responds accurately across the entire evolution for both major and minor species. For the major species the predictions are almost perfect. For lower concentration species, *e.g.,* ethane, the accuracy is lower (see Fig. 8.3(c)), as expected. However, interestingly, it appears that prediction errors are damped over the evolution and the ethane yield at the outlet is in fact well predicted.

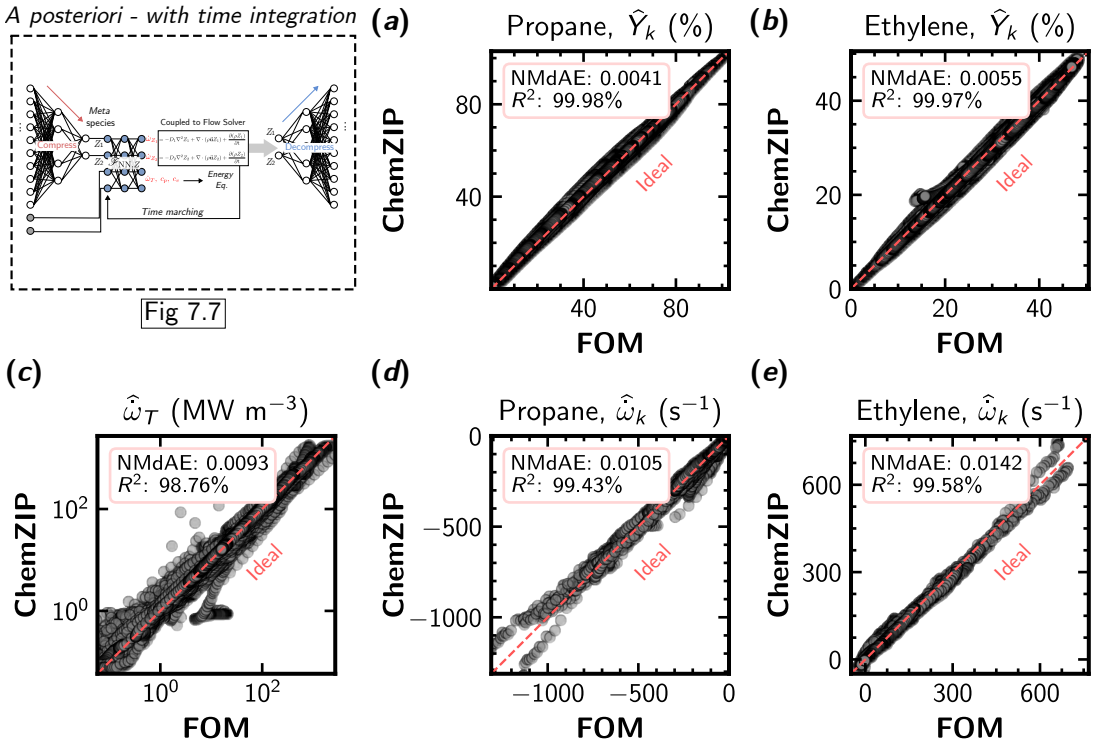


Figure 8.4: A posteriori (*i.e., with* time integration) parity plot (500,000 data points) comparing reconstructed `ChemZIP` predictions against the ground truth (FOM) data for **(*a*)**–**(*b*)** the propane/ethylene dry yields **(*c*)** the heat absorption rate (on a log scale) and **(*f*)**–**(*e*)** the propane/ethylene production rates

To determine whether the conclusions drawn above generalise to all time histories, the complete integrated test database is plotted in Fig. 8.4 and summarised in Table 8.1. The

accuracy of the yields is higher than that of the production rates, which again suggests that errors in the production rates are damped by the integration operator. This is favourable because the yields are the quantities that the reaction system designer is interested in.

Compared with a priori (non-integrated) testing shown in Figs. 8.1 & 8.2, it appears that there is no significant accuracy degradation due to error accumulation. This indicates that the model is robust. A possible explanation for this is the inherent stability and self-correcting nature of an endothermic dynamical system. If the temperature is too high due to modelling errors, the reaction heat absorption increases and brings down the temperature toward the ground truth; conversely, if the temperature is too low, the reaction heat absorption decreases and brings the temperature back up. This physical mechanism suppresses the growth of errors and encourages robustness.

The reliability of `ChemZIP` is confirmed in Fig. 8.5, which shows the distribution of prediction errors over 500,000 test points. Both histograms are centred on a relative error of less than **10%** in all cases with very little change in the error distribution between the a priori and a posteriori tests. However, in both cases, Fig. 8.5 shows a small fraction of relatively high percentage errors ($(x - \widehat{x})/x$). This is primarily due to poor numerical conditioning and bias of this metric as the values of the ground-truth in the denominator

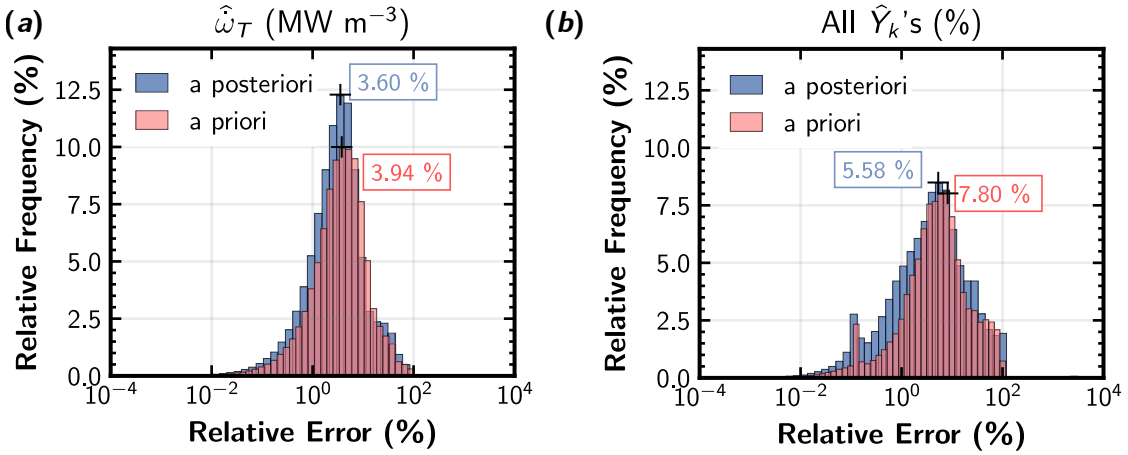


Figure 8.5: Histograms for the relative error $(x) = |(x - \widehat{x})/x|$ between `ChemZIP` and the ground truth FOM for time-integrated (a posteriori) and non-time-integrated (a priori) test data for **(*a*)** the reaction heat absorption rate and **(*b*)** all the species mass fractions. The abscissa is on a log scale

become small. This effect is most pronounced for the mass fractions, which are often precisely equal to zero. However, Fig. 8.4 and other investigations (not shown) indicate

that a poor relative error close to zero does not affect the global solution. Furthermore, focusing on the reaction heat absorption rate as it represents an integrated effect of dynamics in all species, Fig. 8.6 suggests that generally larger relative errors occur in regions of low reaction activity. This is because of limited training data coverage in this low temperature regime. The impact on the yield should be small because the predicted values are still below or close to the activation threshold of the reaction, and the model can be manually turned off in this regime. Despite this, more generally, the analysis

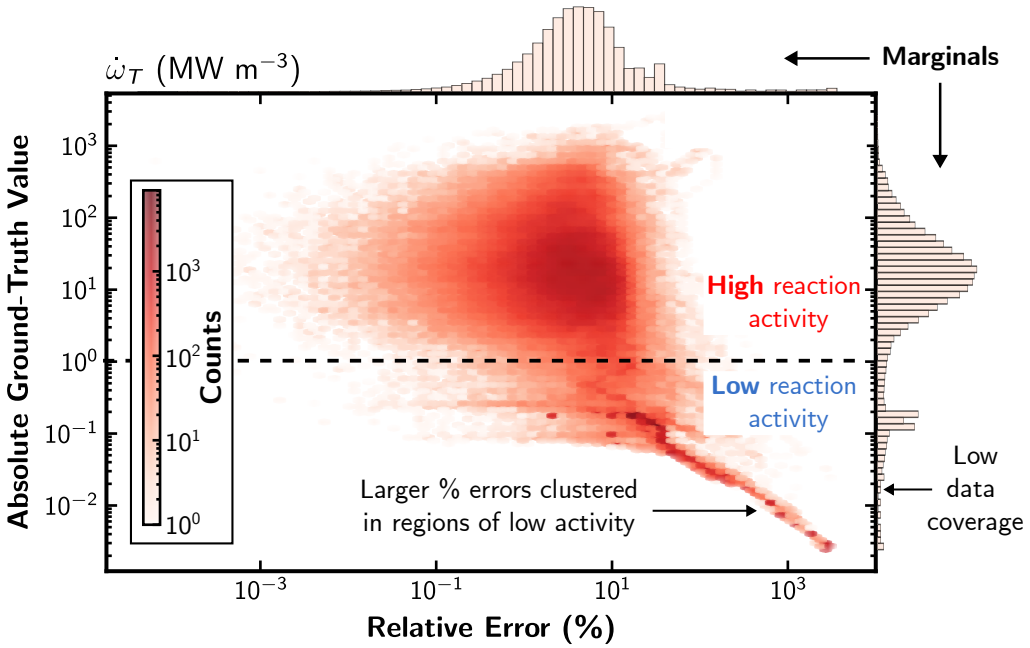


Figure 8.6: 2D joint and marginal distributions for the relative prediction errors and absolute value for the reaction heat absorption rate (500,000 test data points). Both axes are on a log scale

presented provides valuable information that can be fed back into the training data generation process. Given these insights, the modeller can cluster more training samples in under-represented regions to counteract this bias if necessary.

Ensuring consistent accuracy across several orders of magnitude for reaction responses and species yields underscores perhaps the principal challenge in optimising a neural network for chemical kinetics: the inherent multi-scale dynamics of the system. This leads to various optimisation challenges summarised as follows: (1) the optimiser tends to focus on certain scales (2) large variability in scales within a mini-batch can result in difficult learning dynamics and instability in the gradient descent optimisation algorithm and (3) the predictions are highly sensitive to hyperparameter selection, which may be

optimum for some scales but not for others. Some of these challenges are remedied by choosing an appropriate loss function, as described in Sec. 7.3.4; however, further refinement is still needed.

Optimisation Demonstration

This subsection demonstrates that `ChemZIP` can accurately predict trends in quantities of interest under an optimisation-like parameter sweep. A 1D FPR model is used with virtual heat flux patches to represent enthalpy input within a 10-stage turbo-reactor. The level of enthalpy input in each stage is varied over a wide range of realistic conditions from $100 - 450\,\mathrm{kJ\,kg^{-1}}$ (equivalent to $1.0 - 4.5\,\mathrm{MJ\,kg^{-1}}$ of total specific energy supply). Figure 8.7 shows that `ChemZIP` performs well under all enthalpy supply conditions and

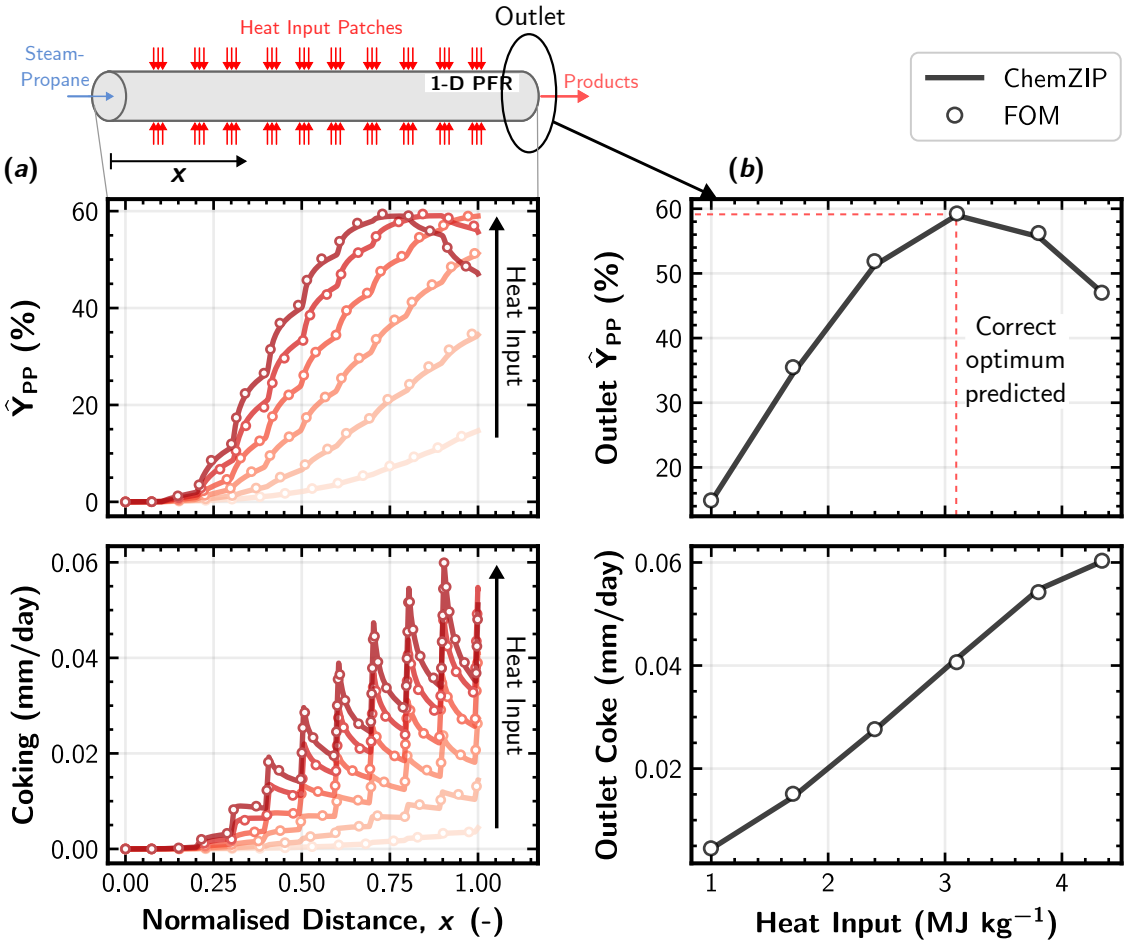


Figure 8.7: Comparing `ChemZIP` and the FOM predictions under parametric design space exploration for the heat input, showing **(*a*)** the high-value primary product (ethylene + propylene + 1 3-butadiene + hydrogen + 1-butene + 2-butene) yield $\mathrm{Y_{PP}}$ and steady state coking rate distributions and **(*b*)** the coking rate and primary product yields at the outlet as a function of heat input

that the optimum primary product yield, as well as the trend in the coking rate, are accurately predicted. Note that the coking rate is calculated during post-processing by first reconstructing the full species mass fraction vector from latent space and then feeding it into the coking model presented in Appendix A.

Although the study presented is somewhat contrived, it is an important verification of model generalisability since the imposed heat flux profiles were not drawn from the same statistical distribution as the data-generating process.

8.3 Verification: Three-Dimensional Heated Duct Flow

To verify the generalisability of `ChemZIP` to more complex multidimensional viscous flow environments, a good test case is a 3-D heated duct. Given that all of the training data was generated from 1D fully-mixed idealised simulations, this test case is important to demonstrate generalisability. A duct was chosen over an axisymmetric pipe because it features 3D effects in the form of corner vortices.

To benchmark the computational performance of `ChemZIP` and verify its predictive accuracy, it is compared against an industry-standard commercial reacting flow solver FLUENT with a detailed kinetic model (*i.e.,* the FOM) as described in Sec. 6.2. For performance benchmarking, `ChemZIP` is also compared against state-of-the-art acceleration strategies available in FLUENT, that is, a combination of DACM + DCZ (see Sec. 6.3.1), which will be denoted as "Fluent Accel." in later figures. Unfortunately, we could not enable ISAT acceleration because there appeared to be an issue with its parallelisation.

8.3.1 Numerical Method and Boundary Conditions

Figure 8.8 shows the computational domain and mesh. The duct dimensions and boundary condition (see Table 8.2) are are representative of a section of a typical steam cracking coil. The heat input $\dot{Q}$ is imposed as a Neumann boundary condition uniformly for all four walls. Employing a long domain (*i.e.,* 5 m) is necessary to ensure the feed undergoes substantial transformation, which can be a slow process due to thermal diffusion limitations of heat transfer. To improve visualisation, contour plots throughout this chapter will be presented using an axially-squashed domain, as shown in Fig. 8.8.

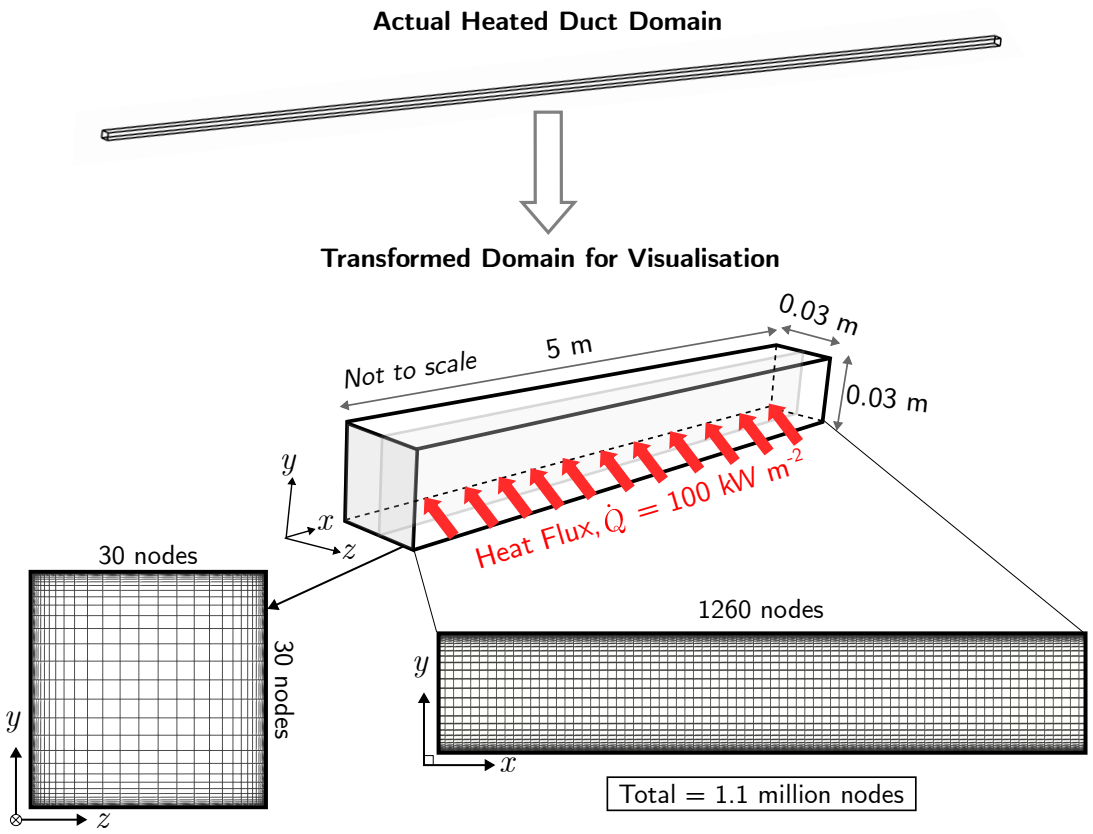


Figure 8.8: Heated duct test case domain and computational mesh (totalling 1.1 million nodes)

The same mesh is used for the simulations in FLUENT and in `ChemZIP`. The maximum non-dimensional wall distance in the domain is $y^+ = 7$; therefore, adaptive velocity wall functions are employed with the standard Spalart-Allmaras turbulence model in both solvers. Since we are not interested in the absolute accuracy of our model, only the relative accuracy compared to the baseline, the quality of the mesh is not so important. Therefore, a finer mesh was avoided because the computational cost of the FOM is already prohibitively costly. Unless stated otherwise, it can be assumed that simulations are performed in parallel on a high-performance Intel Cascade Lake 18-core CPU.

Table 8.2: Boundary conditions for heated duct

Boundary Conditions	
Inlet total pressure (bar)	2.1
Exit static pressure (bar)	2.0
Inlet total temperature (°C)	650
Wall heat flux, $\dot{Q}$ (kW m^{-2})	100 & 150
Inlet turbulent to laminar viscosity ratio	10

The numerics for TBLOCK have been explained in previous chapters. For FLUENT, a steady, compressible, coupled, pressure-based solver is used combined with a second-order finite-volume method. Time integration is performed with an implicit scheme. For most of the solver setup, default settings are used. In both cases, the solutions are fully converged, as illustrated by the propane yield in Fig. 8.9, which is the slowest quantity to reach a steady state. Interestingly, the convergence history for `ChemZIP` features minimal oscillations, indicating that the model is stable.

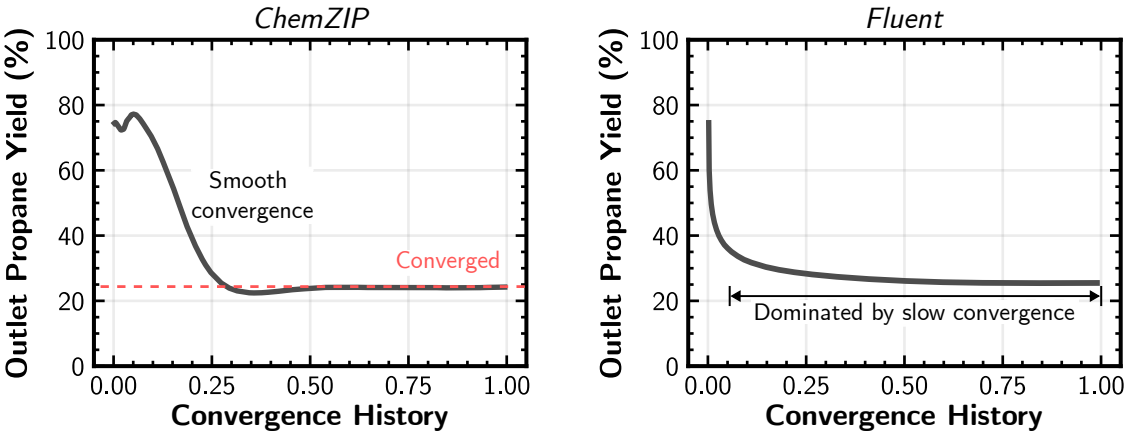


Figure 8.9: Convergence history for propane yield, which is the slowest variable to reach steady state ($\dot{Q} = 150\,\mathrm{kW\,m^{-2}}$)

Baseline Non-Reacting Flow Solver Comparison Check

As explained in Sec. 7.5, the `ChemZIP` module is implemented in the CFD solver TBLOCK, which is optimised to solve flows within turbomachines (which the reader is reminded is our ultimate goal). However, ideally, for verification purposes, `ChemZIP` should also be coupled directly into FLUENT. This is currently not possible due to the complexity of the development and is deferred to a later date. Therefore, this creates a discrepancy when performing verification, as the baseline numerics of FLUENT and TBLOCK are different. For example, pressure-based vs. density-based solver, implicit vs. explicit integrators, different spatial discretisation implementations, different levels of numerical dissipation particularly at low Mach numbers, different wall function implementation, etc. Therefore, the first step is to compare the solution for the non-reacting baseline solvers.

Figure 8.10 indicates that the temperature distributions are similar, which is important because this is the primary driver of the reaction. There is, however, a notable discrepancy in the mass flow rate. This will lead to different residence times. Although the benchmarking and verification results presented below will be somewhat distorted by this systematic bias, the discrepancy is fairly small and, therefore, should not have a significant effect. This systematic difference between solvers should, however, be borne in mind when interpreting the results in the following two subsections.

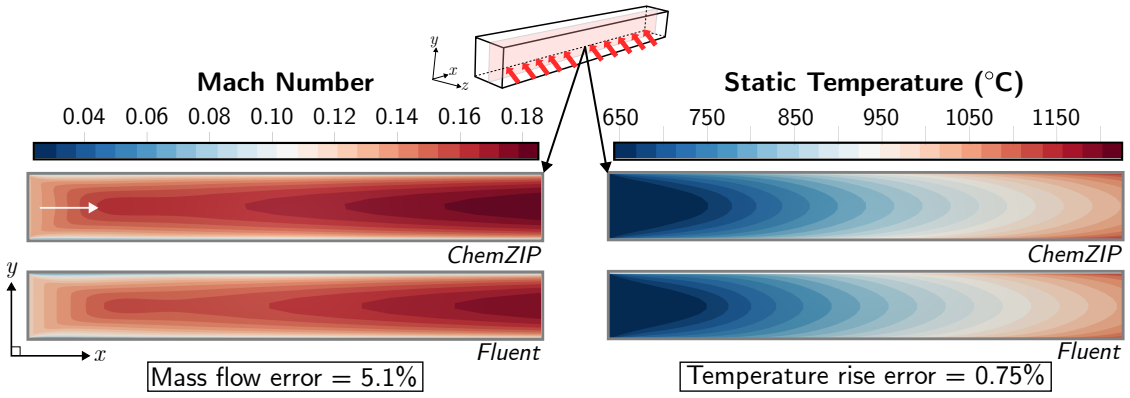


Figure 8.10: A comparison of static temperature and Mach number contours on an mid x-y plane cross section computed from a non-reacting flow simulation ($\dot{Q} = 100\,\mathrm{kW\,m^{-2}}$) for both solvers

8.3.2 Computational Acceleration

Performance benchmarking is achieved by measuring the wall-clock time required to march the solution to a steady state for (a) the FOM in FLUENT (b) an accelerated scheme for the chemistry and (c) `ChemZIP`. Figure 8.11 demonstrates that significant computational acceleration can be achieved using `ChemZIP`. For a 65-species mechanism, a speed-up of **580**× is measured compared to direct integration and **40**× compared to state-of-the-art acceleration strategies. Based on scaling rules, it is estimated that for a more detailed (and more accurate) mechanism, the rewards are even more substantial, with `ChemZIP` potentially capable of at least a **75**× speed-up over "Fluent Accel." (see Fig. 8.11). Furthermore, greater speed-ups are possible for larger mesh sizes.

Although there is a difference in the underlying flow solver, for performance benchmarking, the evaluation of the chemical source terms dominates, and differences in the underlying numerics of the two solvers should not influence performance measurements

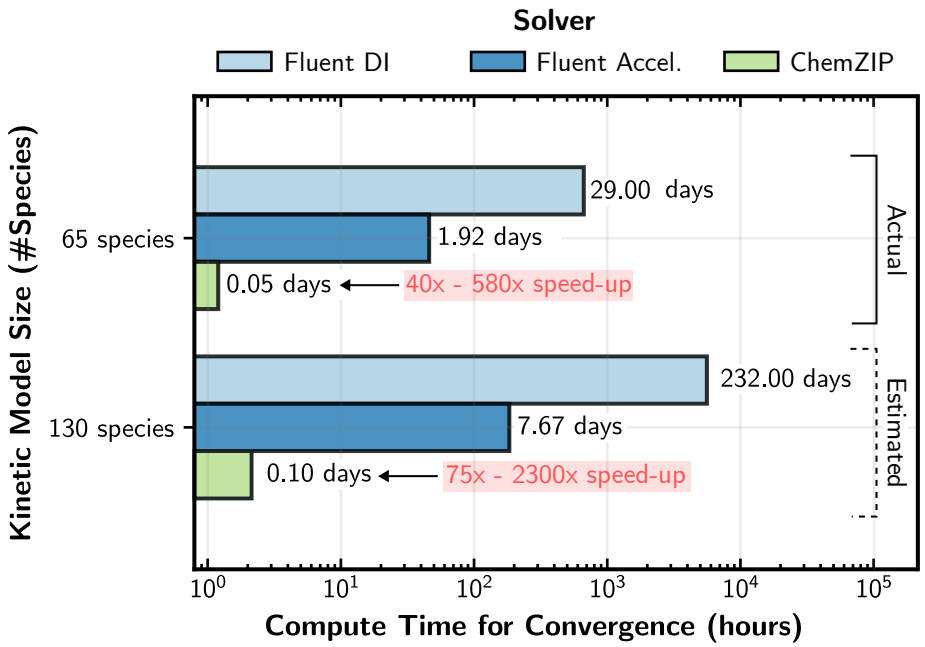


Figure 8.11: Computational acceleration achieved by `ChemZIP` compared with direct integration (DI) in FLUENT and a state-of-the-art acceleration (Accel.) strategy that combines DACM + DCZ

significantly. In fact, for low-Mach-number flows ($Ma < 0.2$) such as the test case presented, it is likely that the speed-up factors would be further improved if `ChemZIP` were integrated as a module in FLUENT. This is because FLUENT's pressure-based method converges 2.5 times faster for low-Mach-number flows than the density-based method available in TBLOCK. Note that the implicit pressure-based method is the only method available for reacting flows in FLUENT.

Breakdown

For DI, Fig. 8.12 shows that the normalised CPU time per iteration is, by far, dominated by evaluating the chemical source terms (*i.e.,* the net species production rates $\dot{\omega}_k$ and reaction heat absorption rate $\dot{\omega}_T$). This overhead is partially alleviated with DACM + DCZ, but this does not tackle the cost of solving a large number of additional species transport equations, which now dominates CPU time (see Fig. 8.12). With `ChemZIP`, the addition of chemistry only increases the CPU time by a factor of around 3 compared to non-reacting CFD (*i.e.,* mass, momentum & energy equations). It is also important to observe that the cost of interfacing TBLOCK with the TENSORFLOW model through the `Python` interface

(as discussed in Sec. 7.5.1), incurs almost negligible overhead. Therefore, the flexibility gained by coupling `Fortran` with a high-level programming language comes at little cost.

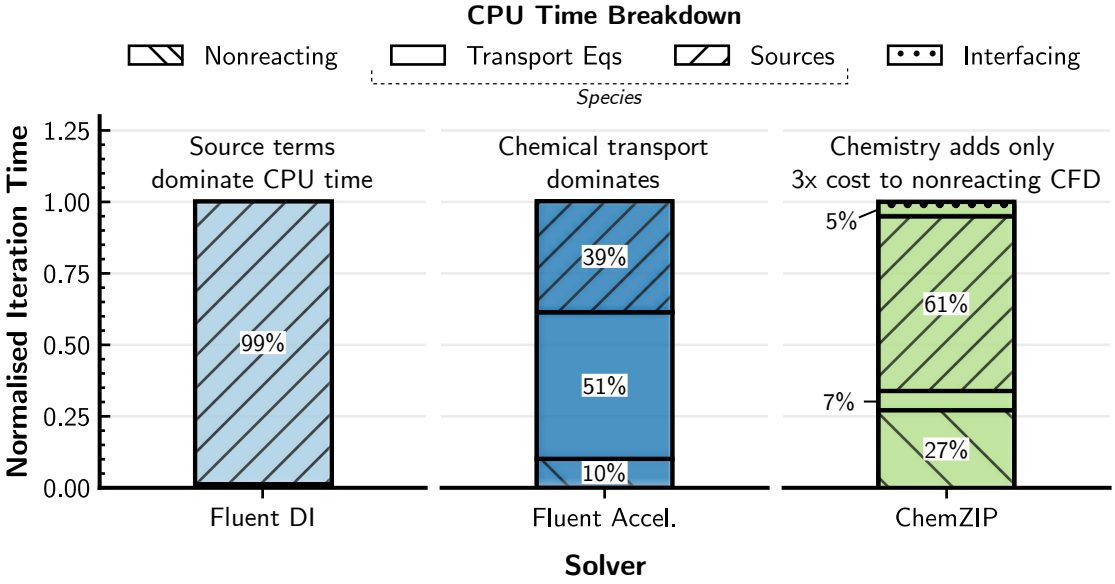


Figure 8.12: Breakdown of the normalised CPU time per CFD iteration, comparing `ChemZIP` with FLUENT for a 65-species mechanism

Figures 8.13 & 8.14 illustrate several of the primary contributions to the computational acceleration achieved by `ChemZIP`. First, the number of transport equations is reduced, in this example, from 65 to 3. Second, the average numerical stiffness of the chemical system is decreased by a factor of approximately 10^9 (see Fig. 8.13(a)). The stiffness is defined as the ratio of maximum to minimum characteristic chemical time-scales $\frac{\max(\boldsymbol{\tau}_{\text{chem.}})}{\min(\boldsymbol{\tau}_{\text{chem.}})}$, where $\boldsymbol{\tau}_{\text{chem.}}$ is calculated from the eigenvalues of the Jacobian matrix $\mathbf{J}$ (see Eq. (6.12)) in physical space or the Jacobian matrix $\mathbf{J}_Z = \frac{\partial \dot{\boldsymbol{\omega}}_Z}{\partial \boldsymbol{Z}}$ in latent space [101]. It is noted that very long time-scales $\tau_{\text{chem.}} > 1000\,\text{s}$ are filtered out, as they are associated with approximately non-reacting species. It is this intense stiffness of the system that necessitates the use of specialised implicit stiff ODE solvers. These require frequent evaluation and inversion of the Jacobian matrix in each mesh cell, effectively, to scale the solution update appropriately for each species. The Jacobian inversion is the reason why the cost of DI grows cubically with the number of species. For the latent dynamical system, fast-reacting intermediate species are filtered out, and thus the stiffness is almost eliminated. Not only does this dramatically improve the rate of convergence to a steady-state solution (see Fig. 8.9), it also enables the use of fast explicit time integrators

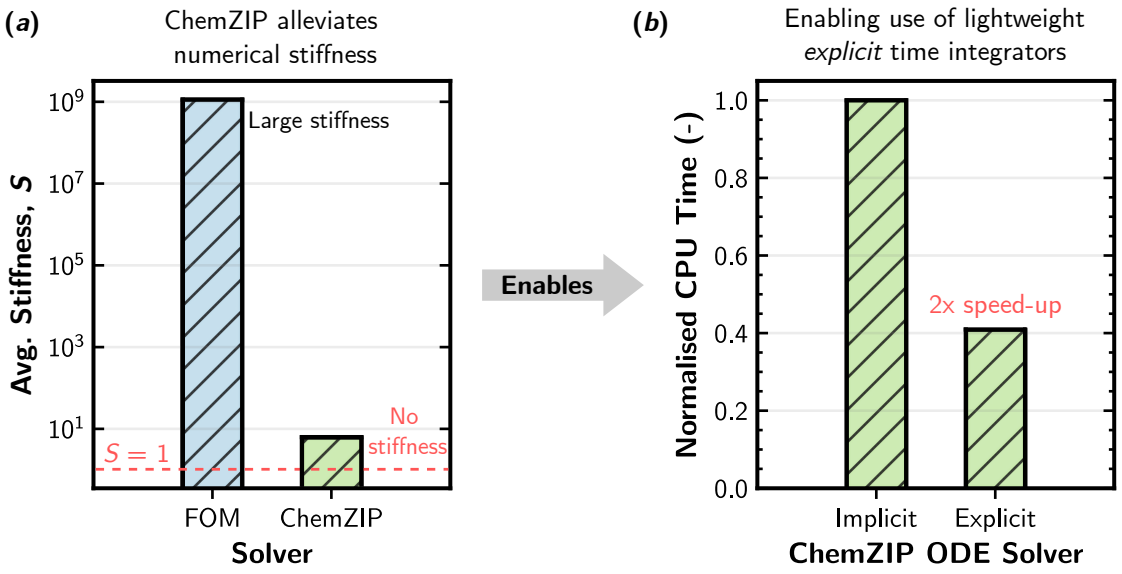


Figure 8.13: For a 65-species propane mechanism: **(*a*)** a comparison of the average numerical stiffness for the chemistry (on a log scale) and **(*b*)** the normalised single-core CPU time required to time integrate the `ChemZIP` model with either an stiff implicit integrator or an explicit integrator

instead. These do not require the Jacobian matrix and incur a much smaller overhead per step. For the same error tolerances, Fig. 8.13(b) demonstrates that a factor of 2 performance benefit is possible by integrating the latent system with an explicit method over a stiff implicit one. Furthermore, Fig. 8.14(a) shows that, unlike the FOM, the latent space system does not impose a restriction on the time-step set by the datum non-reacting flow system. This is because the characteristic chemical time-scale is similar to that of the fluid. Therefore, the baseline numerical methods of the non-reacting CFD solver can be used efficiently and without modification.

Finally, Fig. 8.14(b) illustrates that for more than 100 mesh cells, the cost of evaluating the chemical and energy source terms is $\mathbf{50}\times$ faster using a NN in latent space compared to the FOM. For batches of less than 100 cells, the latency associated with the TensorFlow API overshadows the performance. However, for large batches of queries to the NN, this cost is amortised, and the scaling becomes linear (see Fig. 8.14(b)).

The culmination of these factors results in the level of computational acceleration observed previously in Fig. 8.11. Consequently, for the first time, we have demonstrated that the degree of computational acceleration necessary to bring detailed aerochemical modelling into the design optimisation loop is now possible. The next question is whether

the predictive accuracy is sufficient to perform realistic optimisation.

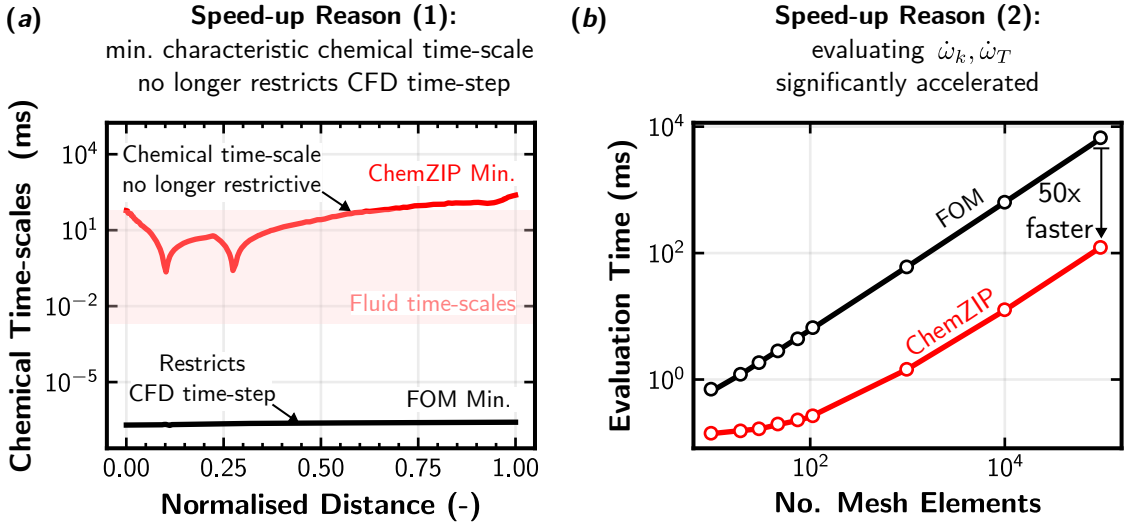


Figure 8.14: For a 65-species propane mechanism: ***(a)*** a comparison of the streamwise variation in minimum characteristic chemical time-scale between `ChemZIP` and the FOM (on a log scale) and ***(b)*** comparison of single-core CPU time to evaluate the chemical and energy source terms as a function of number of mesh cells in a computational batch (on a log-log scale)

8.3.3 Predictive Accuracy

The accuracy of `ChemZIP` is evaluated under two heat flux conditions $\dot{Q} = 100\,\mathrm{kW\,m^{-2}}$ and $\dot{Q} = 150\,\mathrm{kW\,m^{-2}}$ to exercise its performance in a wider region of the thermochemical space. Although this multidimensional viscous flow environment has not been observed during training, Figs. 8.15 & 8.16 show that subject to modelling assumptions (see Sec. 6.2.3), the yields and reaction rates are within several percent of FLUENT. Performance is good for both major species, minor species and temperature in both the streamwise direction and the wall-normal direction, and the contours are surprisingly smooth and symmetric. The channel is long enough (*i.e.,* 5 m) that if compounding of model errors was an issue, it would be observed in the results. This does not appear to be a serious problem, as evidenced by a similar streamwise distribution for the yield and reaction rate between the solvers (see Figs. 8.15 & 8.16). However, due to inconsistencies with the underlying numerics of the two solvers, the streamwise yield is slightly ahead of FLUENT. This is because in TBLOCK the axial velocity is higher (see Fig. 8.10) and the reaction time-scale is lower. For the higher heat flux condition, the subtle stretching of the yield profile is captured by `ChemZIP`, albeit marginally over-stretched.

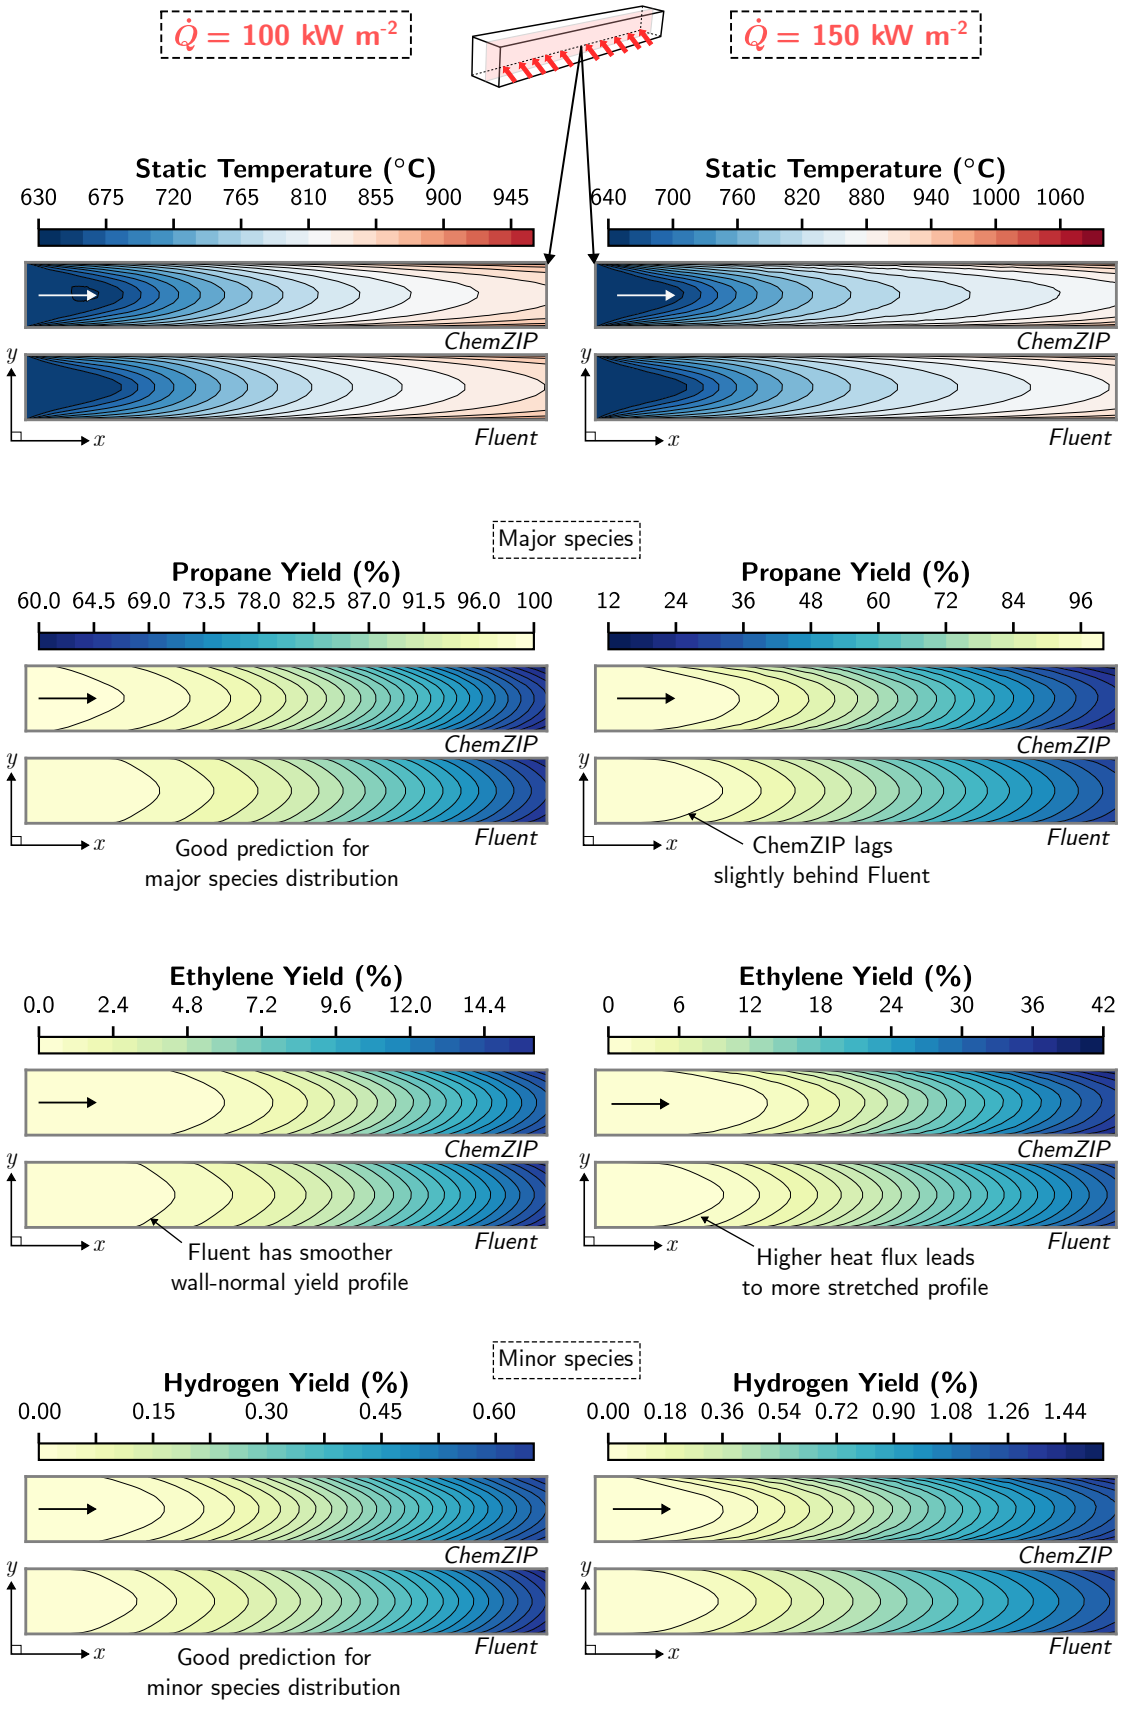


Figure 8.15: Verification of the reconstructed temperature and dry yields under two heat loads

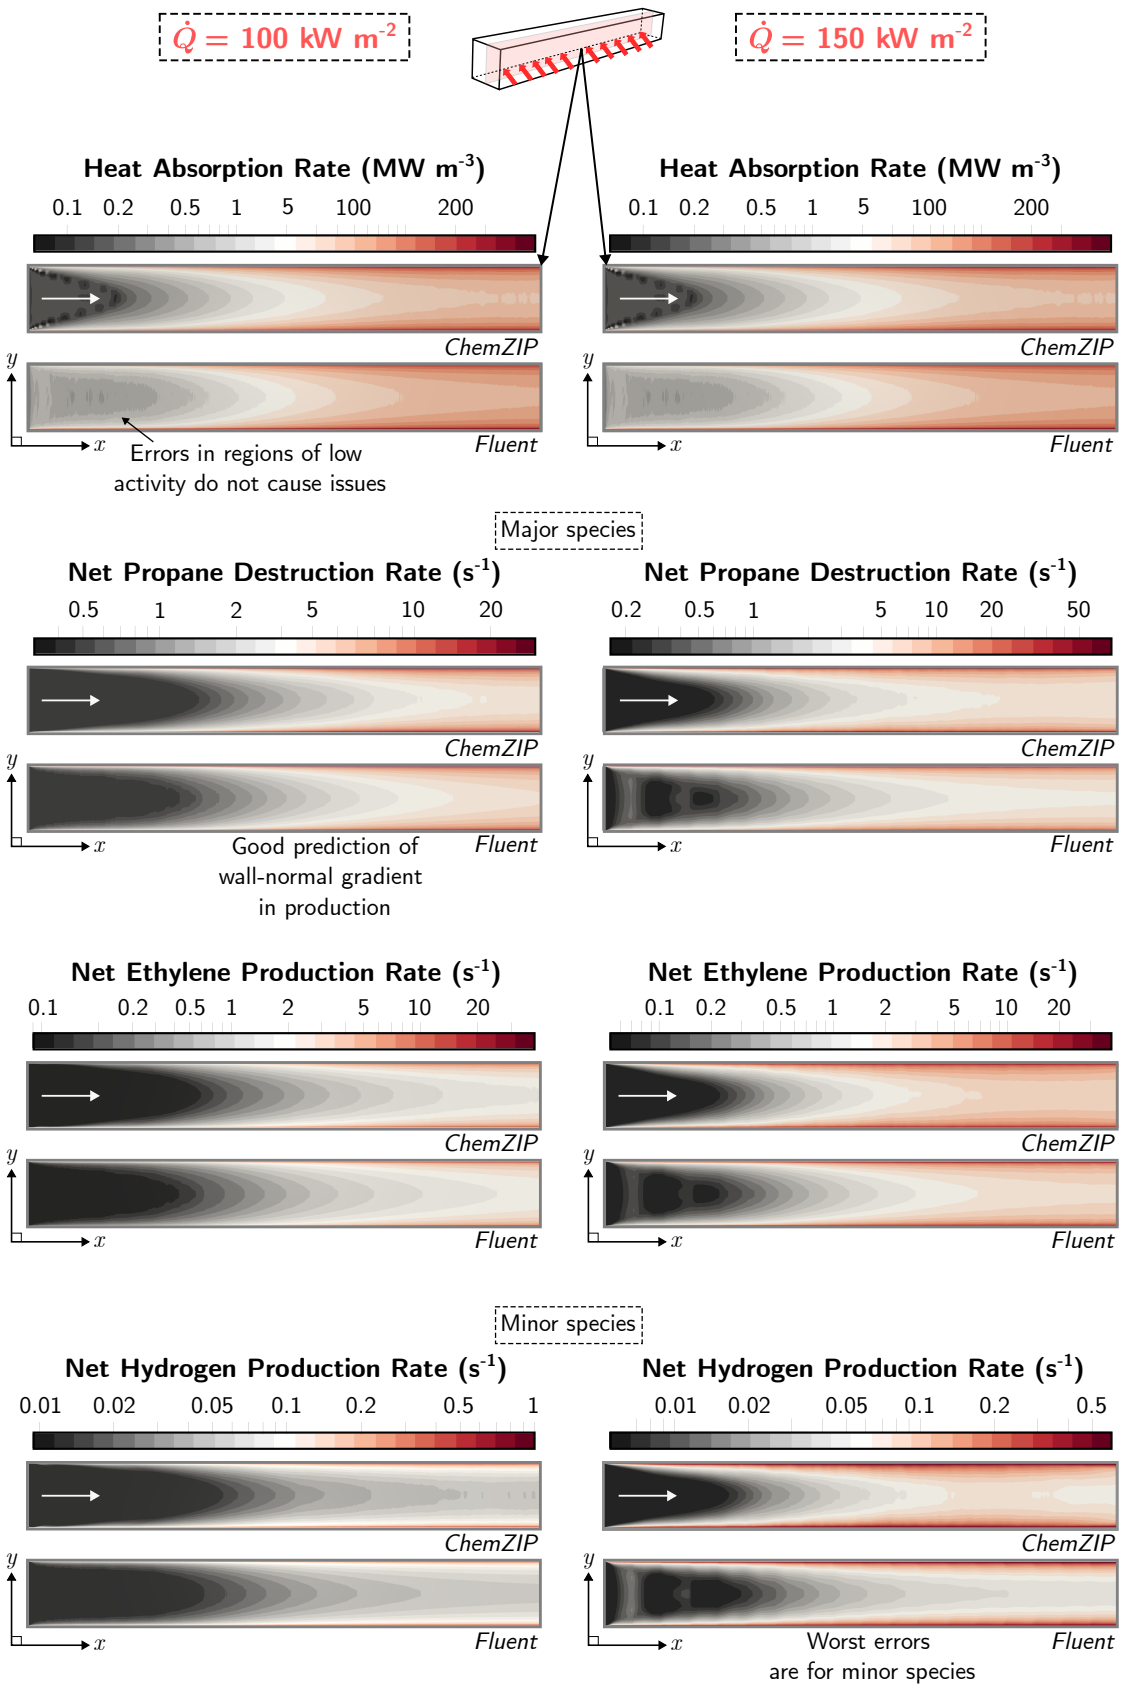


Figure 8.16: Verification of the reconstructed the reaction source terms under two heat loads

Inspecting the yields in Fig. 8.15, there is a small discrepancy in the shape of the wall-normal profile. This difference could be attributed in part to the inconsistency between the underlying numerics of the two solvers. An additional explanation could also be that the reaction system does not respond accurately to turbulent diffusion of the species, which is an approximately temperature-independent process and not incorporated during training. The diffusion term in the FOM produces a gentler wall-normal yield gradient (see Fig. 8.15) due to its smoothing effect. In fact, it is possible that the success of this methodology can perhaps be attributed to the temperature dominance of these types of chemical reaction processes. This is in contrast to other types of reaction processes, such as non-premixed combustion, where diffusion and mixing can become the limiting physical process. On the other hand, while it is true that the reduced-order system may not react perfectly to species diffusion, it is reassuring to see from Fig. 8.17 that diffusion still has a notable effect of the system dynamics despite this mechanism not being included during training data generation.

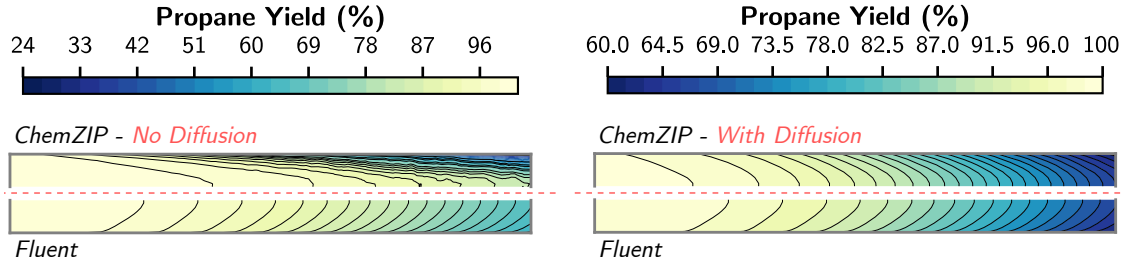


Figure 8.17: A comparison of the propane yield distributions for `ChemZIP` both with and without species diffusion activated ($\dot{Q} = 100\,\mathrm{kW\,m^{-2}}$)

For capturing two- and three-dimensional effects not seen before in this chapter, `ChemZIP` works notably well. Due to lateral temperature gradients, mass diffusion and growth of the fluid boundary layer, a non-uniform reaction rate profile is produced from the core to the wall. In addition, secondary vortices lead to enhanced heat transfer and an accumulation of energy at the corners of the channel (see Fig. 8.18). These 2D/3D effects are mostly well captured by `ChemZIP` (see Figs. 8.16 & 8.18). Remarkably, the reconstructed production rate source terms respond accurately to wall-normal temperature gradients, which is a challenging task due to the high degree of non-linearity and sensitivity of the dynamics. At the inlet, the prediction of the heat absorption rate is poor (see

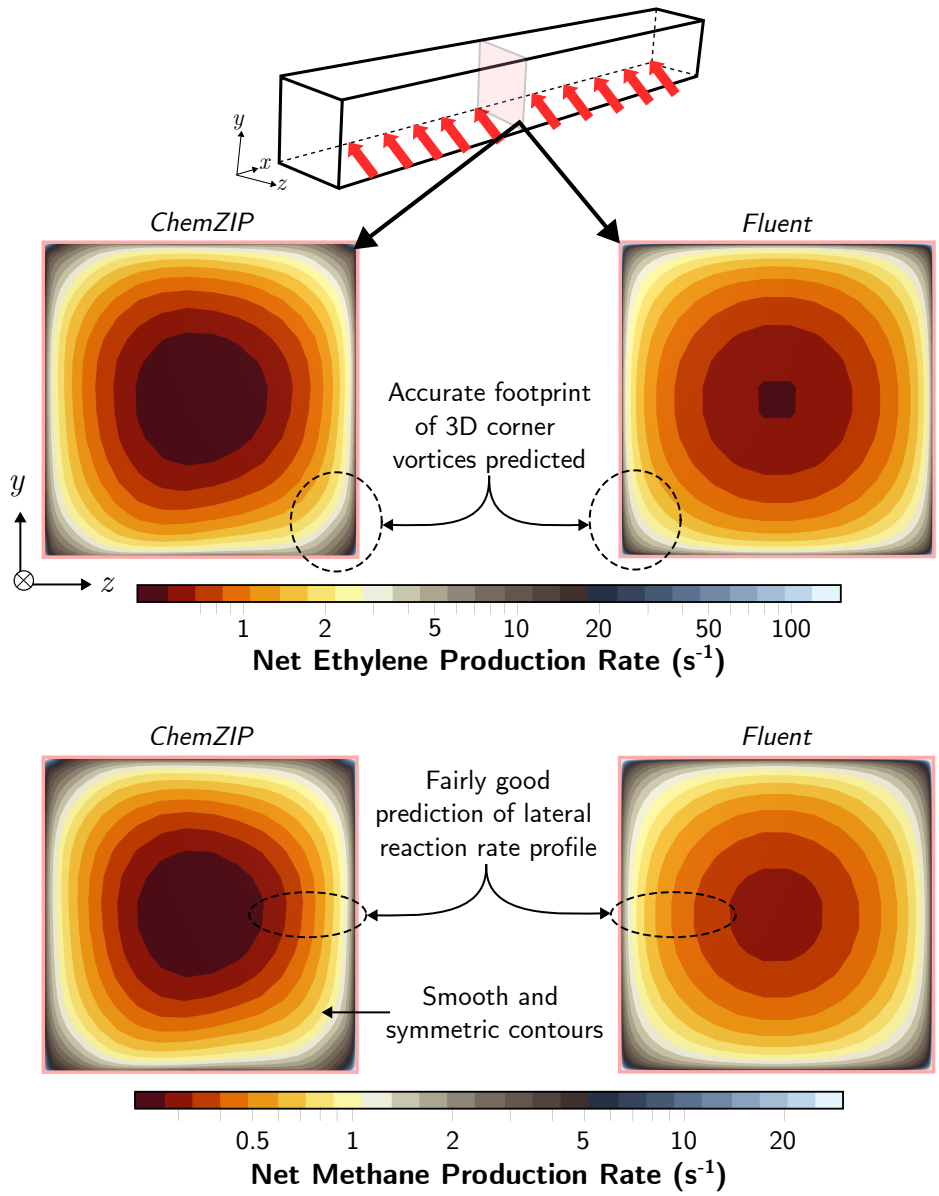


Figure 8.18: Net species production rate distributions on an z-y plane cross section ($\dot{Q} = 100\,\mathrm{kW\,m^{-2}}$)

Fig. 8.16 (Top)) for similar reasons as discussed in Sec. 8.2.2. Again, because the reaction activity is low here, it does not pose a global issue. This is corroborated by the accurate temperature distribution seen in Fig. 8.15 (Top), and recovery of the reconstructed production rates toward the outlet (see Fig. 8.16).

Figure 8.19 shows the latent space CFD solution for all points in the mesh (shown as a red point cloud) superimposed onto the phase space containing the training data manifold. As eluded to previously, due to mass diffusion and two-/three-dimensional effects not captured during training, the regions of the phase space accessed during

the CFD computation are slightly outside the training data manifold. This indicates extrapolation. Nevertheless, we have shown that the accuracy is acceptable and the solution is stable, providing confidence in the generalisability of the model.

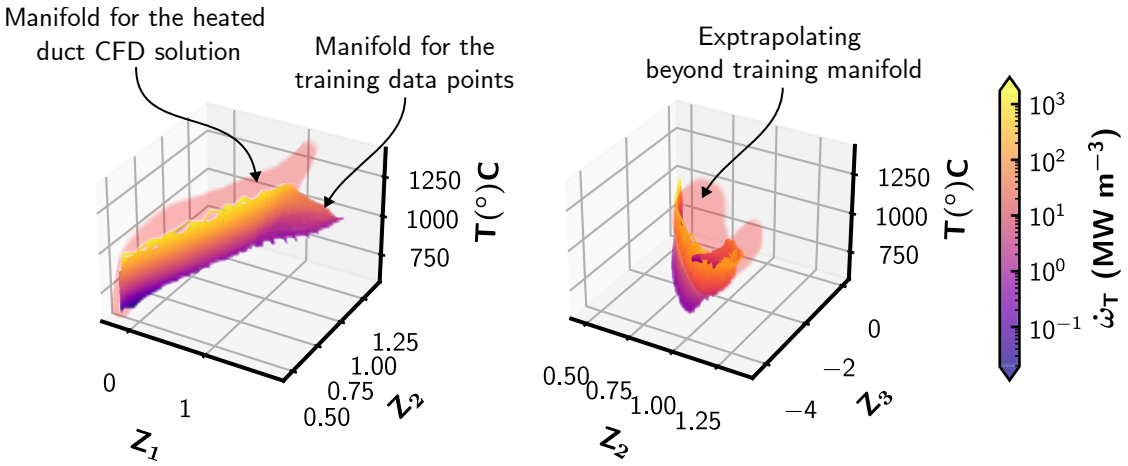


Figure 8.19: A projection of the latent phase space manifolds for all points in the training database and for all solution points from the heated duct ($\dot{Q} = 150\,\mathrm{kW\,m^{-2}}$) CFD simulation (shown in red)

This section has shown that, in addition to the successful performance benchmarking, the accuracy of `ChemZIP` is within reasonable tolerances. Despite the fact that the model was deployed in a very different environment from the 1D training simulator, the spatial profiles of the chemistry were well predicted. This was enabled by intelligent curation of the training database (see Sec. 7.2). Since the chemical source terms are simply an input-output mapping $\mathscr{F}_{\mathrm{NN},Z} : (T, p, \boldsymbol{Z}) \mapsto \dot{\boldsymbol{\omega}}_Z$, training has been performed in such a way as to sample regions of the thermochemical space accessed in more complex flows.

In the following section, `ChemZIP` will be used to solve problems for which it was originally intended: modelling complex aerochemical interactions in the turbo-reactor.

8.4 Proof of Concept: Multi-Stage Turbo-Reactor

As a proof of concept, the new methodology for accelerated aerochemical coupling is now deployed into the turbo-reactor numerical modelling system. The domain is a multi-stage turbo-reactor architecture with an inlet temperature of 840 °C, and the CFD solver is run in transient mode. The rotor stage loading coefficient is approximately $\Psi = 3.0$. This is a uniquely complex aerodynamic environment with rich aerochemical interactions that

were not possible to model before without impractically large computational resources. This provides new insights into the aerochemical flow physics of the turbo-reactor that will inform design optimisation for a more efficient chemical reaction in future iterations.

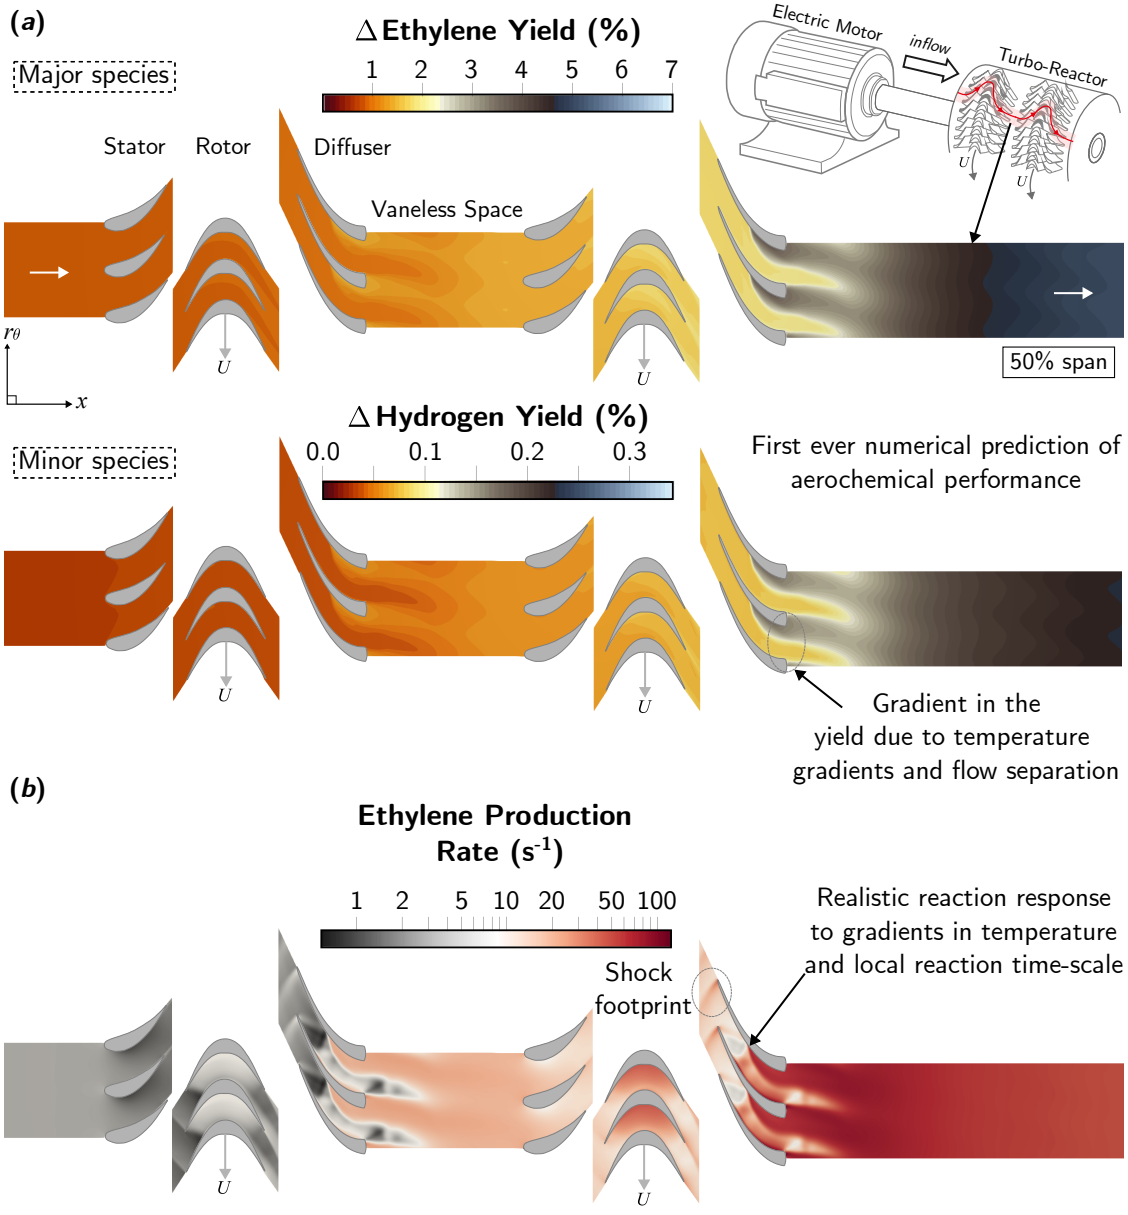


Figure 8.20: Time-averaged **(*a*)** change in major and minor species yields relative to the inlet and **(*b*)** ethylene net production rates (on log scale) at midspan across a two-stage turbo-reactor

Figure 8.20(a) shows the time-averaged yield distribution for both major and minor species yields across two stages. Within the second stage diffuser blade passage, the footprint of separated suction surface flow is clearly visible through a pitchwise gradient in the yields. Therefore, the reaction is responding reasonably to regions with different

reaction time-scales and temperature histories. The ethylene net production rate shown in Fig. 8.20(b) also appears to be producing realistic responses to gradients in temperature. For example, the footprint of several shockwaves (which result in a sharp rise in temperature and hence non-equilibrium response of the reaction) are clearly visible at the diffuser leading edge. Furthermore, by assuming coke formation on the metal surfaces can be approximated solely as a function of the local wall temperature and gas composition in each surface mesh cell, its rate can be extracted from the coking model detailed in Appendix A. Consequently, regions and flow features with a high propensity for coking can be identified early in the design cycle, as illustrated in Fig. 8.21. This could enable considerable improvements in the availability and lifetime of the turbo-reactor.

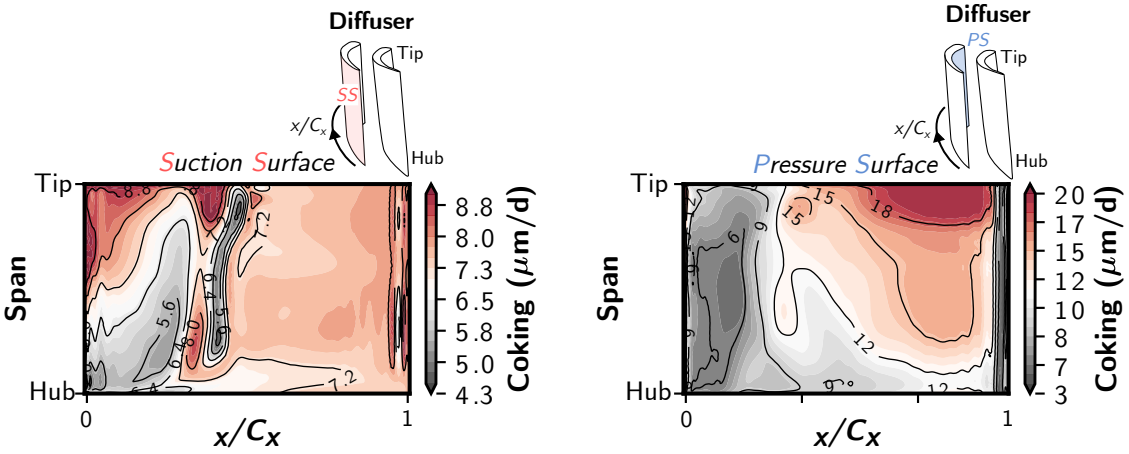


Figure 8.21: Time-averaged steady state coke formation rate distribution in units of µm per day on the suction and pressure surfaces on the diffuser blade in the turbo-reactor

At this point, due to prohibitive computational costs, it is beyond the scope of this work to verify the accuracy of `ChemZIP` for these types of highly complex flows. However, what these results show is that the model provides remarkably realistic and stable solutions despite being deployed in an entirely unfamiliar aerodynamic environment. Once adequate computational resources are available, it will be possible to verify these complex flows.

8.5 Summary

In this chapter, we have comprehensively established the potential of using a machine-learning-enhanced surrogate model to accelerate aerochemical modelling within complex

engineering systems without significantly sacrificing accuracy. A sequence of verification investigations and performance benchmarks has been conducted. Several findings from these experiments can be summarised as follows.

- A posteriori predictions resulted in an R^2 score above 98% for the reconstructed mass fractions and above 95% for the dynamics.
- The a priori and a posteriori error distributions were determined to exhibit minimal differences, suggesting that the system is resilient to error accumulation.
- A key challenge identified during training was the multi-scale nature of reaction kinetics. This complexity led to difficulties in consistently optimising the neural network parameters across a wide range of mass fractions and temperature conditions. This issue will be addressed in future work.
- Performance benchmarking suggested that `ChemZIP` is capable of performing reacting flow calculations up to 75 times faster than state-of-the-art acceleration methods in FLUENT while maintaining errors below 10%. This gave us confidence to deploy `ChemZIP` into the numerical design and modelling system for the turbo-reactor as a proof of concept. A remarkably realistic and stable solution was obtained.
- It has been highlighted that the reduced-order chemistry system may not be accurately responding to species diffusion. The next step could be to incorporate a simplified model for this process into the training database generation system.

So far, for development purposes, the mechanism size and mesh resolution have been limited. For realistic and comprehensive kinetic models of heavy hydrocarbon-based feedstocks used in practical settings, the number of species can grow to $K = 100 - 1000$, and the number of mesh elements could increase several times. This is where conventional acceleration strategies fail. When we apply `ChemZIP` to these problems, it is expected that the level of computational acceleration obtained will increase substantially. This will allow the dynamic response of detailed yet realistic kinetic models to be integrated early into the design cycle.

Chapter 9

Conclusions and Future Work

9.1 Discussion of Major Contributions

This thesis aimed to achieve two main objectives. First, we sought to demonstrate the feasibility of using a novel turbomachinery concept to efficiently decarbonise high-temperature chemical reaction processes (up to 1700 °C). A critical component of this was developing a deeper understanding of the uniquely complex aerothermal interactions within these machines. Second, we sought to address the gap in available numerical tools for design, modelling and analysis of the turbo-reactor, particularly for chemistry and aerochemistry computations. In this section, the main contributions of this work are summarised in four parts corresponding to the research questions laid out in Sec. 1.2.

9.1.1 The New Turbo-Reactor Concept

> *Research question 1: Is the turbo-reactor an efficient and effective solution to impart a large amount of mechanical energy into a working fluid over short length and time scales?*

The new turbo-reactor concept has been proven to be an efficient, effective and financially attractive solution for decarbonising over 40 hard-to-abate industrial sectors. By replacing surface heat exchange within a furnace with direct mechanical energy transfer to the fluid within the turbo-reactor, the power density increases by a factor of 50 to 500, while simultaneously enhancing the quality (*e.g.,* reaction efficiency) of the process. For the first time, high- and low-fidelity numerical computations have been used to confirm the feasibility of the concept, demonstrating that a level of specific energy input 10 to 20 times higher than that of an axial compressor can be imparted stably and at an approximately constant global pressure ratio. Therefore, the energy input required for a wide range of endothermic chemical reactions can be supplied within

a shaft length of less than 2 m. It was confirmed that turbulent mixing, necessary to dissipate mechanical energy between stages, decayed sufficiently rapidly for well-conditioned flow to enter the downstream stages.

9.1.2 Energy Conversion Mechanisms

> *Research question 2: What are the physical mechanisms responsible for the conversion of mechanical energy into internal energy?*

It has been explained that the turbo-reactor brings about a fundamental change in turbomachinery design philosophy. Rather than the designer being concerned with aerodynamic efficiency, their objective is instead to supply the ideal *quantity* of energy to the reaction for its completion and control its *distribution* in order to design a temperature distribution that maximises reaction efficiency. The temperature profile is determined by the balance between reaction heat absorption and rate of energy conversion (proportional to the entropy generation rate), which is composed of various mechanisms. Detailed numerical investigations concluded that the dominant energy conversion mechanism was turbulent mixing generated by the thick diffuser trailing edge. Despite the supersonic Mach number regime ($Ma > 1.2$) the combined reversible and irreversible shock contributions were lower than expected at about 20% of the total, primarily from reversible pressure work. A more important effect was from shock interactions with the boundary layer. Across a wide range of Reynolds numbers, the energy conversion process appeared to be approximately Reynolds number independent. The effect of the Mach number was more significant, with a 20% increase in the total energy conversion coefficient between $Ma = 1.1$ and $Ma = 1.5$; however, the scaling was weaker than expected due to the counterbalancing effects of irreversible and reversible energy conversion contributions.

9.1.3 Robustness and Flexibility of the Turbo-Reactor

> *Research question 3: Is the elemental stage design concept robust & flexible and can a universal stage design philosophy be used for a wide range of working fluids?*

Feed flexibility: In Sec. 4.2 we investigated whether it was possible to use a universal stage design philosophy across a wide range of feedstock compositions (*i.e.,* for ammonia decomposition, hydrogen production, ethane cracking and n-hexane cracking), reaction states and operating conditions. A pragmatic framework was developed to capture the unidirectional effect of the reaction progression on aerothermal stage performance in a decoupled manner without having to run prohibitively costly reacting flow calculations. The turbo-reactor was shown to be robust and flexible across a broad range of conditions, operating stably and only showing a 15% reduction in the stage loading coefficient across all feedstocks and reaction states. Of course, in practice, the stage design will likely need to be tuned stage-by-stage and for different applications/feedstocks. However, the analysis presented suggests that the required modifications will be relatively minor.

Multi-stage design tailoring: A two-stage turbo-reactor was studied with different interstage vaneless space lengths ranging from $1C_x$ to $3C_x$. Within the interval, it was concluded that the performance of the downstream stage is independent of interstage vaneless space length. This gives the aerochemical designer the flexibility to adjust this distance to optimise the temperature profile for improved chemical reaction efficiency. Using idealised 1D chemical kinetic calculations with an ethane feedstock, the multi-stage arrangement was optimised over a series of different parameter spaces. In this simplified modelling environment, for a fixed coking rate, it was found that a 10% relative increase in the primary product yield could be achieved by non-uniformly separating the stages. A further 7% improvement was possible by allowing a variable level of enthalpy input per stage. At an integral level, this idealised analysis has illustrated that there is considerable untapped opportunity to improve the reaction efficiency by tailoring the aerodynamic design. To put this into context, an improvement in the yield of 20% could reduce annual feedstock costs for the industry by up to $30 billion. However, the chemical kinetic simulations are overly simplified and the anticipated improvements may not be realised in practice. Furthermore, we have not yet developed any understanding of the effect of realistic aerochemical interactions and spatial non-uniformities on reaction efficiency and coke formation rate.

9.1.4 Accelerating Aerochemical Computations

Research question 4: Can a multi-fidelity, data-driven methodology be used to accelerate the modelling of aerochemical interactions in the turbo-reactor with sufficiently high accuracy that they can feature routinely in parametric design optimisation and exploration?

Methodology: It has been made clear that aerochemically coupled CFD computations are a necessary requirement to optimise the aerodynamic design for maximum reaction efficiency with low coking rates. This is especially true for the turbo-reactor, which features complex aerochemical interactions absent in conventional reaction tubes. Unfortunately, reacting flow calculations are prohibitively costly due to the high-dimensionality and multi-scale nature of realistic chemical mechanisms. These problems are most pronounced for high-carbon-number, heavy feedstocks used in the hydrocarbon cracking industry. State-of-the-art acceleration techniques like DACM do not provide a sufficiently powerful speed-up for reacting flow computations to be performed at the design level.

Therefore, a new multi-fidelity machine-learning-enhanced framework was developed to address these limitations. A simple and fast methodology for generating training data in only a few hours was developed. The high-dimensional composition vector is compressed onto a low-dimensional basis using a data-driven compression algorithm. The dynamic response projected into latent space is approximated with a neural network mapping. This approach is more universal than FGM since redundancies in the governing physics can be automatically eliminated for any reaction process; and it also may be more accurate by exploiting non-linearity. Neural networks are unrivalled in their ability to scale to ever larger datasets, which is why they are adopted in this work over other methods like random forests or Gaussian processes. Therefore, in future work, the training dataset can be scaled to tens of millions of points to enhance the generalisability of `ChemZIP`.

Using source terms, the surrogate kinetics are coupled to a set of transport equations for the meta-species. Currently, the flow solver is Tblock, but this could also be a 2D throughflow solver for fast stage design purposes; an unstructured GPU-based solver like U-Tblock; or a commercial solver like Fluent if this is more convenient for industry.

Benchmarking and verification: Initial performance benchmarking indicated that `ChemZIP` was 40 times faster than FLUENT when employing state-of-the-art acceleration strategies. However, this was for a small mechanism. For more accurate mechanisms and large grids, it is likely that the measured speed-ups will be $\mathcal{O}(100\times - 1000\times)$. The observed computational acceleration is a result of three factors (1) 20 times fewer transport equations (2) 50 times lower cost evaluating the energy and chemical source terms and (3) 10^9 times reduction in stiffness, enabling application of fast explicit integrators and eliminating the Jacobian evaluation and inversion (which is a bottleneck).

A series of verification studies successfully demonstrated that the accuracy of `ChemZIP` was sufficiently good at this proof-of-concept stage. The R^2 score was above 95% in all cases and the error distributions were very similar for the a priori and a posteriori tests, indicating that the model was remarkably robust to error accumulation. The largest errors were observed to occur in regions of low reaction activity, where the data coverage was low. This does not significantly affect the global solution because these points are close to or below the activation energy. However, more generally, the multi-scale nature of chemical kinetics poses challenges for the learning algorithm that need to be addressed.

When `ChemZIP` was deployed in an entirely unfamiliar 3D viscous flow environment (*i.e.,* a heated 3D duct), the reconstructed yields and production rates were within 10% of FLUENT for a range of different species and levels of heat load. The effects of diffusion and two/three-dimensional non-uniformities were well captured despite being absent during training. However, some enhancements are necessary to improve the prediction of the wall-normal yield profiles. Finally, `ChemZIP` was implemented into the turbo-reactor numerical modelling system. A remarkably realistic and stable solution was obtained. Even at this proof-of-concept stage, these results provide new insights into aerochemical interactions.

This work has successfully demonstrated that a multi-fidelity machine-learning-enhanced approach can be used to accelerate chemical kinetic computations sufficiently that they can be brought into the aerodynamic design optimisation loop of the turbo-reactor. To advance this platform beyond the proof-of-concept stage, there are several recommendations for future work as outlined below.

9.2 Recommendations for Future Work – Aerochemistry

In this section, we offer a set of suggestions for future development of the new aerochemical modelling framework.

9.2.1 Network Architecture and Learning Algorithm

Training: In Chapter 8 it was indicated that inter-species diffusion in a realistic reacting flow environment pulled the latent phase space trajectories away from the training data manifold, leading to extrapolation. Although this was not a significant issue and the response from the source terms in these unexplored regions was surprisingly well captured, it would be desirable to expand our training database. One strategy would be to improve the noise augmentation model discussed in Sec. 7.2.3, potentially by preserving the atomic balances and thus constraining the samples generated to a more appropriate region of the composition space, as discussed by Wan *et al.* [180]. Moreover, noise augmentation should be applied selectively to a limited number of target species rather than indiscriminately to all species. Alternatively, an additional term can be added to the PFR reactor equations used for training data generation [180]. This does not need to be accurate but should simply capture the smoothing effect of diffusion over a range of randomly sampled time-scales.

During the verification process, it was also observed that predictions in low-temperature regions of the thermochemical space were deficient. This was in part due to insufficient training data coverage. This could easily be remedied by artificially adding more time-steps or increasing the sample weightings in these regions. Alternatively, a *mixture of experts* approach (see Almeldein & Van Dam [151]) could be taken in which an ensemble of "expert" sub-networks learn to specialise on sub-regions of the training data space, with a gating network to smoothly transition in-between.

Architectural: Up to this point, the aim of this work has been to demonstrate the feasibility of the `ChemZIP` framework. This was done without concentrating on fine-tuning the learning algorithm and neural network structure, or integrating recent advancements from the machine learning community (see Bishop & Bishop [167]). At the most basic level, there is scope to optimise the architectural hyperparameters, including the number

of hidden layers, number of neurons, number of latent meta-species; incorporating various modern advanced regularisation techniques like dropout, particularly for deeper networks; and/or using batch normalisation to smooth the loss function landscape and learning dynamics. Neural ODEs [181–184] or neural operators [152] could also be explored to approximate the latent space dynamics more robustly.

To balance accuracy across the multiple scales present in reaction dynamics, Döppel & Votsmeier [146] introduced a custom activation function in the last layer of the neural network to improve learning behaviour and demonstrated an eight-fold reduction in relative error. This could help alleviate many of the multi-scale challenges we observed.

Inference: Presently, given the current thermochemical state, the latent space source terms are retrieved by querying a neural network (with over 4000 parameters) in each cell of the mesh. This was shown to account for 61% of the CPU time per iteration. Depending on the size of the NN, it may be computationally cheaper to use the ISAT methodology to tabulate this mapping on the fly. This could work well as the number of meta-species is small $\mathcal{O}(3)$.

9.2.2 Physical Modelling Assumptions

For simplicity, several assumptions about the reacting flow system have been made (see Sec. 6.2.3). Although several of these assumptions are valid for the applications examined in this thesis, this is not universally the case and therefore should be addressed for broader applicability of `ChemZIP`.

Variable molecular diffusivities: It has been assumed that the physical/meta species have a constant molecular mass diffusivity coefficient independent of temperature and composition. This assumption has a negligible impact on the freestream flow, where turbulent diffusion dominates. However, near the wall, this assumption may need to be refined. This can be important when catalysts are present at the surface or for accurate coking rate predictions. To remedy this, the thermochemistry-dependent molecular diffusivities should be stored in the training database. During training, they would be

projected into latent space and a mapping would be learnt from the latent thermochemical state to the laminar diffusion coefficient for each meta-species. Unfortunately, variable-diffusivity modelling will add an additional computational cost for online operation.

Add cross-species heat flux: To reduce computational overheads, the cross-species heat flux term in the energy equation (Eq. (6.4))

$$q_{st} = \nabla \cdot \left(\rho \sum_{k=1}^{K} h_k \bar{D}_k \nabla Y_k \right)$$

has been neglected. Adding this term to the flow solver will be costly. Within each mesh cell, the *low-dimensional* latent space composition will have to be projected back into the *high-dimensional* physical space where q_{st} is calculated. This should only be done if q_{st} is of a magnitude similar to or greater than $\dot{\omega}_T$ (which it is in many cases).

Turbulence–chemistry interaction: The very high level of turbulence in the turbo-reactor could potentially have a notable impact on the reaction dynamics. This is due to TCI, primarily from temperature fluctuations. Modelling this effect is different from typical combustion-based approaches, which are primarily centred around the process of fuel-oxidiser mixing. Instead, consider a periodic temperature fluctuation. Due to the exponential form of the reaction rate (see Eq. (3.1)), temperatures above the mean will increase the reaction rate more than those below the mean. As a result,

$$\overline{\dot{\omega}_k}(T, p, Y_1, \ldots, Y_K) \neq \dot{\omega}_k \left(\overline{T}, \overline{p}, \overline{Y}_1, \ldots, \overline{Y}_K \right).$$

Details of TCI modelling can be found in Refs. [113, 114, 135], but the basic idea is to assume that temperature is a random variable following an assumed probability density function (PDF), such as a Gaussian distribution $T \sim \mathcal{N}(\overline{T}, \sigma_T^2)$. The production rates can then be integrated over this distribution as follows

$$\overline{\dot{\omega}_k} \left(\overline{T}, \overline{p}, \overline{Y}_1, \ldots, \overline{Y}_K, \sigma_T^2 \right) = \int_0^\infty \dot{\omega}_k \left(T, \overline{p}, \overline{Y}_1, \ldots, \overline{Y}_K \right) \cdot \mathcal{N}(\overline{T}, \sigma_T^2) \, dT.$$

The same integration can be done for the latent space source terms. The temperature scalar variance σ_T^2 becomes a transported quantity [113, 114]. The price paid for this additional modelling fidelity is 5 to 10 additional NN evaluations for accurate numerical quadrature.

One way to avoid this is to perform this integration during training for different randomly sampled values of σ_T^2 within a typical range of $0.1\,\mathrm{K} \leq \sigma_T \leq 50\,\mathrm{K}$ [113]. It follows that the latent space source terms would now be approximated by

$$[\overline{\dot{\boldsymbol{\omega}}}_{\boldsymbol{Z}}, \overline{\dot{\omega}}_T, \ldots]^\top = \mathscr{F}_{\mathrm{NN},Z}(\overline{T}, \overline{p}, \overline{\boldsymbol{Z}}, \sigma_T^2).$$

In other words, the NN inputs will now include the temperature variance and the outputs will be the Reynolds-averaged source terms. An alternative approach to incorporate TCI would be Lagrangian Monte Carlo particle-based methods [114].

Potentially, efficient analysis of TCI effects is now possible using `ChemZIP` + a PDF-based method. This could be a major step toward an accurate modelling and design system for the uniquely complex aerochemical interactions in the turbo-reactor.

9.2.3 Aerodynamic–Coke Interaction

It has been shown that the coke formation rate on the metal surfaces can now be estimated assuming that a good coking model is available. Presuming that all coke that forms stays attached to the surface as an immovable deposit (which may not be true), the change in the blade surface profile can be estimated from the rate of coking. At fixed intervals of virtual plant operation, reacting flow calculations (using `ChemZIP`) can be re-run with updated blade profiles to assess the change in aerothermal performance due to coke-induced geometric variability, as illustrated in Fig. 9.1. An example implementation of this dynamic coke layering approach can be found in Zhang *et al.* [185].

9.2.4 Catalytic Reaction Modelling

Many, if not most processes in the chemical industry require catalysts. This includes synthetic fuel production, ammonia cracking, ammonia synthesis, hydrogen production, catalytic converters, and many more. Therefore, an important extension of `ChemZIP`

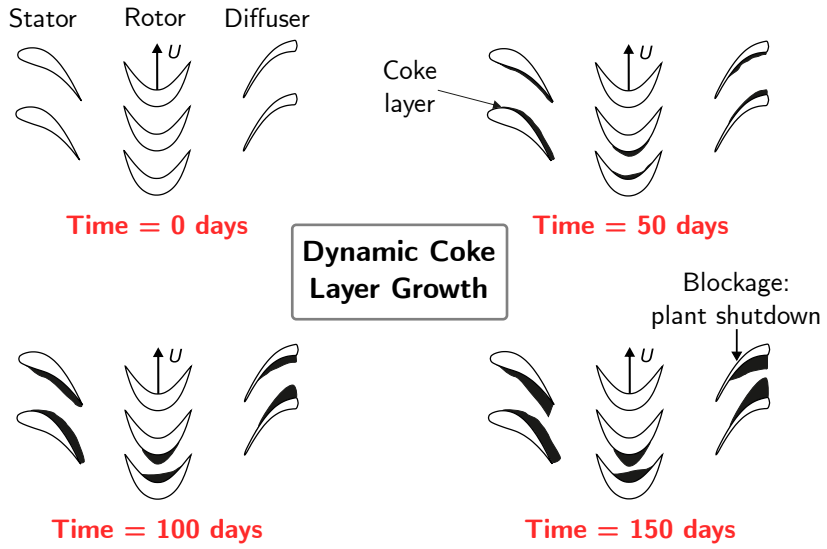


Figure 9.1: Illustration of coke growth on the blade surfaces

is the modelling of surface-phase reactions. This will require two separate models, one for the gas phase and one for active sites[†] on the surface phase. The evolution of surface sites is modelled by an ODE

$$\frac{d\theta_k}{dt} = \frac{\dot{s}_k}{\Gamma}, \tag{9.1}$$

where θ_k are the site fractions, $\dot{s}_k$ are the net production rates of surface species, and Γ is the total density of surface sites. In each surface mesh cell, this equation can be integrated in time.

At the gas–surface interface, fluxes are matched through a set of Neumann boundary conditions. For example, the wall boundary condition for the gas-phase species is determined by balancing the diffusive flux with the net production rate of surfaces species, as shown below

$$\nabla \cdot \left(\left(\rho \bar{D}_k + \frac{\mu_t}{Sc_t} \right) \nabla Y_k \right) \Big|_{\text{wall}} = M_k \dot{s}_k. \tag{9.2}$$

There is also a corresponding flux-matching condition for the energy equation. More details can be found in Kee *et al.* [90].

[†]these are locations on the catalyst surface into which a gas-phase species can be adsorbed, undergo the catalysed reaction, before deadsorbing

9.2.5 Future Applications

Beyond the catalytic applications mentioned above, there are numerous potential areas of multiphysics modelling into which `ChemZIP` could be extended. For any multiphysics problem in which the modelling of the high-dimensional, stiff, multi-scale chemical kinetics is the computational bottleneck, this framework could apply. In many of these cases, additional physics will need to be implemented into the framework. For example, for non-equilibrium high-temperature hypersonic flows [186], an additional transport equation is needed to account for energy exchanges with vibrational degrees of freedom [102]. In the context of nuclear fusion and plasma-assisted combustion/processes, this methodology could be extended to plasma reactions by incorporating electronic impact and vibrational energy transfer reactions [187]. These are already integrated in Cantera, which is currently used for training data generation. Similarly, to model climate and atmospheric chemistry [188], photochemically activated reactions must be included in the training database. This also applies to even more exotic applications, such as modelling chemistry in the atmosphere of an exoplanet, as shown by Hendrix *et al.* [189].

In summary, `ChemZIP` could be used to accelerate multiphysics modelling in a variety of applications ranging from combustion, hypersonics and catalytic chemical processes (*e.g.,* hydrogen production) to atmospheric modelling, electrochemistry (batteries and electrolysers) and plasma dynamics.

9.3 Final Remarks

This step-change advancement in the field of computational science brings multiphysics modelling from being a "one-off" simulation to being exploited routinely at the design optimisation level. Whilst historically, this development cycle would require many physical experiments, now, by integrating complex multiphysics early in the design cycle, the burden of experiments can be reduced and costs saved, accelerating the net-zero transition.

Appendices

Appendix A

Coke Formation Model

The semi-empirical coke formation model used in this thesis is based on the combined catalytic and pyrolytic coking model developed by Zhang *et al.* [190]. The coke formation rate is assumed to be a function of the local temperature, pressure, and concentrations of precursor species at the gas-coke interface. This model is based on the notion that coke is formed through a series of reactions between gas-phase precursor species and active sites on the surface of the metal or on the coke itself. Catalytic coking is a heterogeneous catalytic reaction between metal particles, such as nickel, and the gas-phase precursors. The rate rapidly decays to zero within a day or so once all active metal sites have been consumed. Pyrolytic coking is a non-catalytic reaction between gas-phase precursors and the coke layer formed prior through the catalytic mechanism. The complete list of $N_{\text{precus.}}$ precursor species includes ethylene, propylene, 1-butene, butadiene, cyclopentadiene, benzene, toluene, styrene and xylene. One or more of these precursors is adsorbed into surface sites and undergoes dehydrogenation to form carbon and hydrogen. For example, for an ethylene (C_2H_4) precursor species, the pyrolytic coking reaction is as follows:

$$\mathrm{nC^{\bullet}} + \mathrm{C_2H_4} \longrightarrow (\mathrm{n}+2)\,\mathrm{C^{\bullet}} + 2\,\mathrm{H_2} \tag{A.1}$$

where $\mathrm{C^{\bullet}}$ represents a coke site. Both catalytic and pyrolytic coking are represented by Arrhenius rate equations. The former is modelled by the following equation:

$$r_{\text{coke,cata.}} = \frac{dC_{\text{cata.}}}{dt} = \underbrace{\exp\left(\frac{-\delta_c}{L_c}\right)}_{\text{decays to zero}} \sum_{i}^{N_{\text{precus.}}} A'_{\mathrm{C}_n\mathrm{H}_m} \exp\left(\frac{-E_{a,\text{cata.}}}{RT}\right) C_{\mathrm{C}_n\mathrm{H}_m} \tag{A.2}$$

where $r_{\text{coke,cata.}}$ is the catalytic coking rate, $C_{\text{cata.}}$ is the catalytic coke concentration, δ_c is the coke layer thickness, L_c is the critical thickness, $A'_{\mathrm{C}_n\mathrm{H}_m}$ is the pre-exponential factor for each precursor species, $E_{a,\text{cata.}}$ is the activation energy and $C_{\mathrm{C}_n\mathrm{H}_m}$ is the local concentration of a precursor species.

Pyrolytic coking can be modelled with a reaction equation of a form similar to Eq. A.2 except without the leading exponential decay term. For both models, the empirical coefficients—including the activation energies, rate constants, pre-exponential factors, etc.—are regressed from experimental data, as shown in Ref. [190]. The database was created by measuring coking at laboratory scale for two different feedstocks within the temperature range 800 °C to 950 °C and feed conversion range 35% to 75%.

However, it should be noted that both the ranges of temperature and feed conversion are set by typical conditions in conventional furnaces. Since the turbo-reactor operates slightly outside these conditions, there is uncertainty in the validity of the coking model in this regime. Furthermore, a clear distinction should be made between coke formation, which is what this model is aimed at, and coke deposition. The quantity of coke formed may not be the same as that deposited. This is because some carbon may spall off due to shear stresses or differential thermal expansion. However, the coking rate in the model presented here is dependent solely on the local static temperature at the gas-coke interface and concentrations of the precursor species.

Appendix B

Detailed Neural Network Architecture

B.1 Detailed Network Structure

Tables B.1–B.3 shows details of the neural network architecture for the encoder, decoder and latent dynamics shown previously in Fig. 7.7 in Chapter 7. This setup was developed for a 65-species propane pyrolysis chemical kinetic model generated automatically using RMG-Py. It is emphasised that the number of network parameters is relatively low because low-latency inference is the primary objective.

Table B.1: Detailed neural network architecture for the encoder

No. of layers†	No. Neurons	Activations
2	$21 \rightarrow 3$	Linear

† this includes the input layer (=21× after cull) and bottleneck layer (=3×)

Table B.2: Detailed neural network architecture for the decoder

No. of layers†	No. Neurons	Activations
6	$3 \rightarrow 4 \rightarrow 7 \rightarrow 10 \rightarrow 20 \rightarrow 21$	ELU (4×), linear (1×)

† this includes the bottleneck layer (=3×) and output layer (=21× after cull)

Table B.3: Detailed neural network architecture for the dynamics mapping

No. of layers	No. Neurons	Activations
7	$3 \rightarrow 48 \rightarrow 32 \rightarrow 28 \rightarrow 24 \rightarrow 20 \rightarrow 21$	GELU (5×), sigmoid (1×)

The total number of trainable network parameters is 66 for the encoder, 263 for the decoder and 4039 for the dynamics. Currently, the dimensionality of the latent space bottleneck layer is set to 3 as a trade-off between speed and accuracy; however, this can be set to any value. The non-linearities (*i.e.,* activation functions) are exponential linear

units (ELU) and Gaussian linear error units (GELU), which have been shown to speed-up training and improve accuracy [191]. However, in the output (*i.e.,* last) layer for the NN shown in Table B.3, a *linear* activation function is used to avoid restricting the dynamic range of the output. For the output of the decoder, a *sigmoid* non-linearity is used to bound the mass fractions between 0 and 1 according to the definition of mass fractions.

B.2 Details of the Training Procedure

The stochastic gradient descent method `AdamW` is used for optimisation [192], and regularisation is universally added using weight decay to mitigate overfitting. The backpropagated gradients through the network are clipped to prevent instability. Several of the hyperparameter values are shown in Tables B.4 & B.5, some of which are optimised using KerasTuner. Typically, training takes 20.4 hours on an NVIDIA Quadro RTX 6000 GPU. The training hyperparameters used are shown in Table B.4. The training dataset is split 80%:10%:10% between training, validation and test sets respectively, and the data is randomly shuffled at each training epoch. `EarlyStopping()` to used to automatically stop training once the R^2 score for the validation set no longer improves over a window of 10 epochs. The weights associated with the lowest R^2 score in the moving window are then restored.

Table B.4: Hyperparameters of the learning algorithm

Batch Size	**Learning Rate**	**Weight Decay**	**Epochs**
32	0.0004	0.0004	41

Table B.5: Loss function weights from Eq. (7.7)

λ_1	λ_2	λ_3	w_1	w_2
100	7	0.025	1.0	0.1

References

[1] A. Henry, R. Prasher, and A. Majumdar. "Five Thermal Energy Grand Challenges for Decarbonization". In: *Nature Energy* 5.9 (2020), pp. 635–637.

[2] G. P. Thiel and A. K. Stark. "To Decarbonize Industry, We Must Decarbonize Heat". In: *Joule* 5.3 (2021), pp. 531–550.

[3] Z. Yu. "Grand Challenges in Heat Decarbonisation". In: *Front. Therm. Eng.* 2 (2022), p. 940072.

[4] IEA. "Renewable Energy for Industry". In: IEA Paris, France. 2017.

[5] B. Zühlsdorf et al. "Analysis of Technologies and Potentials for Heat Pump-Based Process Heat Supply Above 150 °C". In: *Energy Conversion and Management: X* 2 (2019), p. 100011.

[6] E. Guillot et al. "Some Details About the Third Rejuvenation of the 1000 KWth Solar Furnace in Odeillo: Extreme Performance Heliostats". In: *AIP Conf. Proc.* 2033.1 (Nov. 2018), p. 040016.

[7] H. Ritchie. "Sector by Sector: Where Do Global Greenhouse Gas Emissions Come From?" In: *Our World in Data* (2020).

[8] IEA. "End-Use Prices Data Explore". In: IEA Paris, France. 2023.

[9] O. Roelofsen et al. "Plugging in: What Electrification Can Do for Industry". In: *McKinsey: New York, NY, USA* (2020).

[10] U. E. I. Association et al. "Levelized Costs of New Generation Resources in the Annual Energy Outlook 2013". In: *US Department of Energy, January* (2013).

[11] D. of Energy and C. Change. *Guidance on Estimating Carbon Values Beyond 2050: An Interim Approach*. 2011.

[12] C. Westbrook et al. *The Role of Comprehensive Detailed Chemical Kinetic Reaction Mechanisms in Combustion Research*. Tech. rep. Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), 2008.

[13] IEA. "Simplified Levelised Cost of Petrochemicals for Selected Feedstocks and Regions". In: IEA Paris, France. 2017.

[14] Coolbrook Oy. *Coolbrook Oy*. Helsinki, Finland, accessed Oct. 15, 2022, https://coolbrook.com. 2022. URL: `https://coolbrook.com` (visited on 02/22/2024).

[15] IEA. "Tracking Clean Energy Progress 2023". In: IEA Paris, France. 2023.

[16] I. Amghizar et al. "Sustainable Innovations in Steam Cracking: CO2 Neutral Olefin Production". In: *React. Chem. Eng.* 5 (2 2020), pp. 239–257.

[17] X. Liu, A. Elgowainy, and M. Wang. "Life Cycle Energy Use and Greenhouse Gas Emissions of Ammonia Production From Renewable Resources and Industrial By-Products". In: *Green Chem.* 22 (17 2020), pp. 5751–5761.

[18] A. Reznichenko and A. Harlin. "Next Generation of Polyolefin Plastics: Improving Sustainability with Existing and Novel Feedstock Base". In: *SN Applied Sciences* 4.4 (2022), p. 108.

[19] H. Eerens and D. van Dam. "Decarbonisation Options for Large Volume Organic Chemicals Production, Dow Chemical Terneuzen". In: (2022).

[20] M. S. Abbas-Abadi et al. "Challenges and Opportunities of Light Olefin Production Via Thermal and Catalytic Pyrolysis of End-Of-Life Polyolefins: Towards Full Recyclability". In: *Prog. Energy Combust. Sci.* 96 (2023), p. 101046.

[21] O. Mynko et al. "Electrification of Steam Cracking As a Pathway to Reduce the Impact of the Petrochemical Industry on Climate Change". In: *J. Cleaner Prod.* 427 (2023), p. 139208.

[22] S. Sadrameli. "Thermal/Catalytic Cracking of Hydrocarbons for the Production of Olefins: A State-Of-The-Art Review I: Thermal Cracking Review". In: *Fuel* 140 (2015), pp. 102–115.

[23] M. W. M. Van Goethem. "Next Generation Steam Cracking Reactor Concept". In: (2010).

[24] M. Fakhroleslam and S. M. Sadrameli. "Thermal/Catalytic Cracking of Hydrocarbons for the Production of Olefins; a State-Of-The-Art Review III: Process Modeling and Simulation". In: *Fuel* 252 (2019), pp. 553–566.

[25] V. Spallina et al. "Techno-Economic Assessment of Different Routes for Olefins Production Through the Oxidative Coupling of Methane (OCM): Advances in Benchmark Technologies". In: *Energy Convers. Manage.* 154 (2017), pp. 244–261.

[26] M. Geerts et al. "Steam Cracking Coke Properties and Their Influence on Furnace Run Length Predictions: Experimental and Modeling Study". In: *Ind. Eng. Chem. Res.* 59.52 (2020), pp. 22460–22472.

[27] A. Ochoa et al. "Coke Formation and Deactivation During Catalytic Reforming of Biomass and Waste Pyrolysis Products: A Review". In: *Renewable Sustainable Energy Rev.* 119 (2020), p. 109600.

[28] J. W. Snoeck, G. F. Froment, and M. Fowles. "Filamentous Carbon Formation and Gasification: Thermodynamics, Driving Force, Nucleation, and Steady-State Growth". In: *J. Catal.* 169.1 (July 1997).

[29] A. E. Muñoz Gandarillas et al. "Influence of the Reactor Material Composition on Coke Formation During Ethane Steam Cracking". In: *Ind. Eng. Chem. Res.* 53.15 (2014), pp. 6358–6371.

[30] J. Towfighi, M. Sadrameli, and A. Niaei. "Coke Formation Mechanisms and Coke Inhibiting Methods in Pyrolysis Furnaces". In: *J. Chem. Eng. Jpn.* 35.10 (2002), pp. 923–937.

[31] J. L. Otegui, J. D. Bona, and P. G. Fazzini. "Effect of Coking in Massive Failure of Tubes in an Ethylene Cracking Furnace". In: *Eng. Fail. Anal.* 48 (2015), pp. 201–209.

[32] S. H. Symoens et al. "State-Of-The-Art of Coke Formation During Steam Cracking: Anti-Coking Surface Technologies". In: *Ind. Eng. Chem. Res.* 57.48 (2018), pp. 16117–16136.

[33] H. Zimmermann and R. Walzl. "Ethylene". In: *Ullmann's Encyclopedia of Industrial Chemistry*. John Wiley & Sons, Ltd, 2009.

[34] G. Heynderickx, E. Schools, and G. Marin. "Simulation of the Decoking of an Ethane Cracker with a Steam/air Mixture". In: *Chem. Eng. Sci.* 61.6 (2006), pp. 1779–1789.

[35] Z. Zhang and L. F. Albright. "Pretreatments of Coils to Minimize Coke Formation in Ethylene Furnaces". In: *Ind. Eng. Chem. Res.* 49.4 (2010), pp. 1991–1994.

[36] S. Biswas et al. "A Review on Synthesis of Methane As a Pathway for Renewable Energy Storage with a Focus on Solid Oxide Electrolytic Cell-Based Processes". In: *Front. Energy Res.* 8 (2020), p. 570112.

[37] A. Kiani et al. "Liquefied Synthetic Methane From Ambient CO2 and Renewable H2 - A Technoeconomic Study". In: *J. Nat. Gas Sci. Eng.* 94 (2021), p. 104079.

[38] M. S. Ghafoori et al. "Techno-Economic and Sensitivity Analysis of Biomethane Production Via Landfill Biogas Upgrading and Power-To-Gas Technology". In: *Energy* 239 (2022), p. 122086.

[39] O. A. Aworanti et al. "Enhancing and Upgrading Biogas and Biomethane Production in Anaerobic Digestion: a Comprehensive Review". In: *Front. Energy Res.* 11 (2023).

[40] J. A. Martins et al. "Novel Adsorption–Reaction Process for Biomethane Purification/Production and Renewable Energy Storage". In: *ACS Sustainable Chem. Eng.* 10.24 (2022), pp. 7833–7851.

[41] IEA. "Outlook for Biogas and Biomethane: Prospects for Organic Growth". In: IEA Paris, France. 2020.

[42] G. D. Oreggioni et al. "Techno-Economic Analysis of Bio-Methane Production From Agriculture and Food Industry Waste". In: *Energy Procedia* 123 (2017). Proceedings of 1st International Conference on Sustainable Energy and Resource Use in Food Chains including Symposium on Heat Recovery and Efficient Conversion and Utilisation of Waste Heat ICSEF 2017, 19-20 April 2017, Windsor UK, pp. 81–88.

[43] I. IRENA. "Green Hydrogen Cost Reduction: Scaling Up Electrolysers to Meet the 1.5 C Climate Goal". In: *Abu Dhabi: International Renewable Energy Agency* (2020).

[44] A. Hodges et al. "A High-Performance Capillary-Fed Electrolysis Cell Promises More Cost-Competitive Renewable Hydrogen". In: *Nat. Commun.* 13.1 (2022), p. 1304.

[45] T. Lundblad, M. Taljegard, and F. Johnsson. "Centralized and Decentralized Electrolysis-Based Hydrogen Supply Systems for Road Transportation - A Modeling Study of Current and Future Costs". In: *Int. J. Hydrogen Energy* 48.12 (2023), pp. 4830–4844.

[46] D. DeSantis et al. "Cost of Long-Distance Energy Transmission by Different Carriers". In: *IScience* 24.12 (2021), p. 103495.

[47] E. Giacomazzi et al. "Hydrogen Combustion: Features and Barriers to Its Exploitation in the Energy Transition". In: *Energies* 16.20 (2023), p. 7174.

[48] A. Haj Ayed et al. "Improvement Study for the Dry-Low-NOx Hydrogen Micromix Combustion Technology". In: *Propuls. Power Res.* 4.3 (2015), pp. 132–140.

[49] O. Mynko et al. "Reducing CO2 Emissions of Existing Ethylene Plants: Evaluation of Different Revamp Strategies to Reduce Global CO2 Emission by 100 Million Tonnes". In: *J. Cleaner Prod.* 362 (2022), p. 132127.

[50] R. Pletka et al. "Capital Costs for Transmission and Substations: Updated Recommendations for WECC Transmission Expansion Planning". In: *Black and Veatch PROJECT* 181374 (2014).

[51] L. Zheng, M. Ambrosetti, and E. Tronconi. "Joule-Heated Catalytic Reactors Toward Decarbonization and Process Intensification: A Review". In: *ACS Eng. Au* 0.0 (0), null.

[52] L. Engineering. *Linde E-Furnace*. 2021. URL: `https://www.linde-engineering.com/en/process-plants/petrochemical-plants/linde-e-furnace/index.html` (visited on 02/15/2024).

[53] BASF. *BASF E-Furnace*. 2022. URL: `https://www.basf.com/global/en/who-we-are/sustainability/whats-new/sustainability-news/2022/basf-sabic-and-linde-start-construction-of-the-worlds-first-demonstration-plant-for-large-scale-electrically-heated-steam-cracker-furnaces.htmll` (visited on 02/15/2024).

[54] L. Zheng et al. "Direct Electrification of Rh/Al2O3 Washcoated SiSiC Foams for Methane Steam Reforming: An Experimental and Modelling Study". In: *Int. J. Hydrogen Energy* 48.39 (2023), pp. 14681–14696.

[55] A. Badakhsh et al. "A Compact Catalytic Foam Reactor for Decomposition of Ammonia by the Joule-Heating Mechanism". In: *Chem. Eng. J.* 426 (2021), p. 130802.

[56] S. T. Wismann et al. "Electrified Methane Reforming: A Compact Approach to Greener Industrial Hydrogen Production". In: *Science* 364.6442 (2019), pp. 756–759.

[57] I. Dincer. "1.7 Energy and Exergy Efficiencies". In: *Comprehensive Energy Systems*. Ed. by I. Dincer. Oxford: Elsevier, 2018, pp. 265–339.

[58] L. Zhou et al. "Investigation of a Novel Porous Anodic Alumina Plate for Methane Steam Reforming: Hydrothermal Stability, Electrical Heating Possibility and Reforming Reactivity". In: *Int. J. Hydrogen Energy* 34.2 (2009), pp. 844–858.

[59] S. Wismann. "Electrically Heated Steam Methane Reforming". English. PhD thesis. Technical University of Denmark, 2019.

[60] S. T. Wismann et al. "Electrified Methane Reforming: Understanding the Dynamic Interplay". In: *Ind. Eng. Chem. Res.* 58.51 (2019), pp. 23380–23388.

[61] M. E. H. Tijani, H. Zondag, and Y. Van Delft. "Review of Electric Cracking of Hydrocarbons". In: *ACS Sustainable Chem. Eng.* 10.49 (2022), pp. 16070–16089.

[62] C. H. Lin et al. "Electrified Thermochemical Reaction Systems with High-Frequency Metamaterial Reactors". In: *Joule* 8.10 (2024), pp. 2938–2949.

[63] V. A. Bolotov et al. "Pyrolysis of Heavy Hydrocarbons Under Microwave Heating of Catalysts and Adsorbents". In: *J. Microwave Power Electromagn. Energy* 46.1 (2012), pp. 39–46.

[64] L. Zheng et al. "Electrified CO2 Valorization Driven by Direct Joule Heating of Catalytic Cellular Substrates". In: *Chem. Eng. J.* 466 (2023), p. 143154.

[65] M. R. Almind et al. "Improving Performance of Induction-Heated Steam Methane Reforming". In: *Catal. Today* 342 (2020). SI: Syngas Convention 3, pp. 13–20.

[66] T. Razzaq and C. O. Kappe. "On the Energy Efficiency of Microwave-Assisted Organic Reactions". In: *ChemSusChem* 1.1-2 (2008), pp. 123–132.

[67] A. P. Harvey et al. "Opportunities for Process Intensification Using Non-Thermal Plasmas". In: *Current Opinion in Chemical Engineering* 41 (2023), p. 100930.

[68] E. Thimsen. "Beyond Equilibrium Thermodynamics in the Low Temperature Plasma Processor". In: *J. Vac. Sci. Technol., B* 36.4 (June 2018), p. 048501.

[69] P. Mehta et al. "Catalysis Enabled by Plasma Activation of Strong Chemical Bonds: A Review". In: *ACS Energy Lett.* 4.5 (2019), pp. 1115–1133.

[70] D. S. Mallapragada et al. "Decarbonization of the Chemical Industry Through Electrification: Barriers and Opportunities". In: *Joule* 7.1 (2023), pp. 23–41.

[71] D. Ding et al. "A Novel Low-Thermal-Budget Approach for the Co-Production of Ethylene and Hydrogen Via the Electrochemical Non-Oxidative Deprotonation of Ethane". In: *Energy Environ. Sci.* 11 (7 2018), pp. 1710–1716.

[72] J. Allen, S. Panquet, and A. Bastiani. "Electrochemical Ammonia: Power to Ammonia Ratio and Balance of Plant Requirements for Two Different Electrolysis Approaches". In: *Front. Chem. Eng.* 3 (2021).

[73] IEA. "World Energy Outlook 2022". In: IEA Paris, France. 2022.

[74] G. Bilicic and S. Scroggins. "2023 Levelized Cost Of Energy+". In: *Lazard, April* 12 (2023).

[75] R. T. Shafiee and D. P. Schrag. "Carbon Abatement Costs of Green Hydrogen Across End-Use Sectors". In: *Joule* (2024).

[76] Z. E. Network. *Renewable Energy Resource Assessment Map*. 2010. URL: `https://zoinet.org/` (visited on 02/22/2024).

[77] J. F. Klapproth, G. N. Ullman, and E. R. Tysl. *Performance of an Impulse-Type Supersonic Compressor with Stators*. Tech. rep. 1952.

[78] J. Piechna. "A Review of Shock Wave Compression Rotary Engine Projects, Investigations and Prospects". In: *Energies* 15.24 (2022).

[79] G. Paniagua et al. "Design and Analysis of Pioneering High Supersonic Axial Turbines". In: *Int. J. Mech. Sci.* 89 (2014), pp. 65–77.

[80] A. Hertzberg, A. T. Mattick, and D. A. Russell. "Apparatus for Initiating Pyrolysis Using a Shock Wave". US Patent 5,219,530. 1993.

[81] C. Knowlen et al. "Petrochemical Pyrolysis with Shock Waves". In: *33rd Aerospace Sciences Meeting and Exhibit*.

[82] L. Nill and T. Jahrmarkt. "Design of a Supersonic Steam Tunnel for Use As a Shock Wave Reactor (Cracking Hydrocarbons)". In: *33rd Aerospace Sciences Meeting and Exhibit*.

[83] H. Abraham, H. S. Glick, and S. William. "Method and Apparatus for Carrying Out Gas Phase Reactions Which Require a High Temperature to Promote the Reaction and Rapid Cooling to Preserve the Reaction Product". US Patent 2,832,665. 1958.

[84] R. K. Masse. "Fluid Dynamics of the Shock Wave Reactor". English. PhD thesis. University of Washington, 2001.

[85] A. T. Mattick et al. "Shock-Controlled Chemical Processing". In: *Shock Waves @ Marseille II*. Ed. by R. Brun and L. Z. Dumitrescu. Berlin, Heidelberg: Springer Berlin Heidelberg, 1995, pp. 209–214.

[86] T. Ahonen et al. "Effect of Electric Motor Efficiency and Sizing on the Energy Efficiency in Pumping Systems". In: *2016 18th European Conference on Power Electronics and Applications (EPE'16 ECCE Europe)*. 2016, pp. 1–9.

[87] C. W. Gao et al. "Reaction Mechanism Generator: Automatic Construction of Chemical Kinetic Mechanisms". In: *Comput. Phys. Commun.* 203 (2016), pp. 212–225.

[88] J. Aerssens et al. "The Merit of Pressure Dependent Kinetic Modelling in Steam Cracking". In: *Faraday Discuss.* 238 (0 2022), pp. 491–511.

[89] D. G. Goodwin et al. *Cantera: An Object-Oriented Software Toolkit for Chemical Kinetics, Thermodynamics, and Transport Processes*. `https://www.cantera.org`. Version 3.0.0. 2023.

[90] R. J. Kee et al. *Chemically Reacting Flow: Theory, Modeling, and Simulation*. 2nd ed. John Wiley & Sons, 2017.

[91] J. Blank and K. Deb. "Pymoo: Multi-Objective Optimization in Python". In: *IEEE Access* 8 (2020), pp. 89497–89509.

[92] M. Ghashghaee and R. Karimzadeh. "Multivariable Optimization of Thermal Cracking Severity". In: *Chem. Eng. Res. Des.* 89.7 (2011), pp. 1067–1077.

[93] H. Seada and K. Deb. "A Unified Evolutionary Optimization Procedure for Single, Multiple, and Many Objectives". In: *IEEE Trans. Evol. Comput.* 20.3 (2016), pp. 358–369.

[94] C. C. Centre. *AramcoMech 3.0 (2018)*. 2018. URL: `https://www.universityofgalway.ie/combustionchemistrycentre/mechanismdownloads/#` (visited on 03/24/2024).

[95] E. Ranzi et al. "Reduced Kinetic Schemes of Complex Reaction Systems: Fossil and Biomass-Derived Transportation Fuels". In: *Int. J. Chem. Kinet.* 46.9 (2014), pp. 512–542.

[96] E. Ranzi et al. "New Reaction Classes in the Kinetic Modeling of Low Temperature Oxidation of N-Alkanes". In: *Combust. Flame* 162.5 (2015), pp. 1679–1691.

[97] C. K. Westbrook et al. "A Comprehensive Detailed Chemical Kinetic Reaction Mechanism for Combustion of N-Alkane Hydrocarbons From N-Octane to N-Hexadecane". In: *Combust. Flame* 156.1 (2009), pp. 181–199.

[98] R. Schaller. "Moore's Law: Past, Present and Future". In: *IEEE Spectr.* 34.6 (1997), pp. 52–59.

[99] T. Lu and C. K. Law. "Toward Accommodating Realistic Fuel Chemistry in Large-Scale Computations". In: *Prog. Energy Combust. Sci.* 35.2 (2009), pp. 192–215.

[100] E. S. Oran and J. P. Boris. *Numerical Simulation of Reactive Flow*. 2nd ed. Cambridge University Press, 2000.

[101] M. B. Eva-Maria Wartha and M. Harasek. "Characteristic Chemical Time Scales for Reactive Flow Modeling". In: *Combust. Sci. Technol.* 193.16 (2021), pp. 2807–2832.

[102] J. D. Anderson Jr. *Modern Compressible Flow with Historical Perspective*. 4th ed. McGraw-Hill, 2020.

[103] A. Frassoldati, T. Faravelli, and E. Ranzi. "Kinetic Modeling of the Interactions Between NO and Hydrocarbons at High Temperature". In: *Combust. Flame* 135.1 (2003), pp. 97–112.

[104] K. K. Kuo. *Principles of Combustion*. 2nd ed. New Jersey: John Wiley & Sons, Inc, 2005.

[105] T. Poinsot and D. Veynante. *Theoretical and Numerical Combustion*. 2nd ed. Edwards, 2005.

[106] K. E. Niemeyer, N. J. Curtis, and C.-J. Sung. "PyJac: Analytical Jacobian Generator for Chemical Kinetics". In: *Comput. Phys. Commun.* 215 (2017), pp. 188–203.

[107] C. F. Curtiss and J. O. Hirschfelder. "Integration of Stiff Equations". In: *Proc. Natl. Acad. Sci. U.S.A.* 38.3 (1952), pp. 235–243.

[108] S.-T. Yu et al. "Basic Equations of Chemically Reactive Flows for Computational Fluid Dynamics". In: *36th AIAA Aerospace Sciences Meeting and Exhibit*.

[109] W. Vincenti and C. Kruger. *Introduction to Physical Gas Dynamics*. Wiley, 1965.

[110] A. Ern and V. Giovangigli. *Multicomponent Transport Algorithms*. Vol. 24. Springer Science & Business Media, 1994.

[111] N. Burali et al. "Assessment of the Constant Non-Unity Lewis Number Assumption in Chemically-Reacting Flows". In: *Combust. Theor. Model.* 20.4 (2016), pp. 632–657.

[112] S. N. Appukuttan. "Large Eddy Simulation of the Steam Cracking Process in Refinery Furnaces". PhD thesis. Institut National Polytechnique de Toulouse, 2022.

[113] P. A. Reyniers et al. "CFD Simulations of Industrial Steam Cracking Reactors: Turbulence–Chemistry Interaction and Dynamic Zoning". In: *Ind. Eng. Chem. Res.* 56.51 (2017), pp. 14959–14971.

[114] R. O. Fox. *Computational Models for Turbulent Reacting Flows*. 2nd ed. Cambridge Series in Chemical Engineering. Cambridge University Press, 2003.

[115] R. Bielawski et al. "Highly-Scalable GPU-Accelerated Compressible Reacting Flow Solver for Modeling High-Speed Flows". In: *Comput. Fluids* 265 (2023), p. 105972.

[116] T. Dijkmans et al. "GPU Based Simulation of Reactive Mixtures with Detailed Chemistry in Combination with Tabulation and an Analytical Jacobian". In: *Comput. Chem. Eng.* 71 (2014), pp. 521–531.

[117] M. J. McNenly, R. A. Whitesides, and D. L. Flowers. "Faster Solvers for Large Kinetic Mechanisms Using Adaptive Preconditioners". In: *Proc. Combust. Inst.* 35.1 (2015), pp. 581–587.

[118] A. S. Walker, R. L. Speth, and K. E. Niemeyer. "Generalized Preconditioning for Accelerating Simulations with Large Kinetic Models". In: *Proc. Combust. Inst.* 39.4 (2023), pp. 5395–5403.

[119] K. Seshadri and N. Peters. "Asymptotic Structure and Extinction of Methane-Air Diffusion Flames". In: *Combust. Flame* 73.1 (1988), pp. 23–44.

[120] T. Lu and C. K. Law. "A Directed Relation Graph Method for Mechanism Reduction". In: *Proc. Combust. Inst.* 30.1 (2005), pp. 1333–1341.

[121] P. Pepiot-Desjardins and H. Pitsch. "An Efficient Error-Propagation-Based Reduction Method for Large Chemical Kinetic Mechanisms". In: *Combust. Flame* 154.1 (2008), pp. 67–81.

[122] P. O. Mestas, P. Clayton, and K. E. Niemeyer. "PyMARS: Automatically Reducing Chemical Kinetic Models in Python". In: *J. Open Source Softw.* 4.41 (2019), p. 1543.

[123] T. Lu and C. K. Law. "Linear Time Reduction of Large Kinetic Mechanisms with Directed Relation Graph: N-Heptane and Iso-Octane". In: *Combust. Flame* 144.1 (2006), pp. 24–36.

[124] L. Liang, J. G. Stevens, and J. T. Farrell. "A Dynamic Adaptive Chemistry Scheme for Reactive Flow Computations". In: *Proc. Combust. Inst.* 32.1 (2009), pp. 527–534.

[125] S. Pope. "Computationally Efficient Implementation of Combustion Chemistry Using in Situ Adaptive Tabulation". In: *Combust. Theor. Model.* 1.1 (1997), pp. 41–63.

[126] L. Lu and S. B. Pope. "An Improved Algorithm for in Situ Adaptive Tabulation". In: *J. Comput. Phys.* 228.2 (2009), pp. 361–386.

[127] S. B. P. M. A. Singer and H. N. Najm. "Operator-Splitting with ISAT to Model Reacting Flow with Detailed Chemistry". In: *Combust. Theor. Model.* 10.2 (2006), pp. 199–217.

[128] Z. Ren et al. "The Use of Dynamic Adaptive Chemistry and Tabulation in Reactive Flow Simulations". In: *Combust. Flame.* 161.1 (2014), pp. 127–137.

[129] J. G. S. Long Liang and J. T. Farrell. "A Dynamic Multi-Zone Partitioning Scheme for Solving Detailed Chemical Kinetics in Reactive Flow Computations". In: *Combust. Sci. Technol.* 181.11 (2009), pp. 1345–1371.

[130] U. Maas and S. Pope. "Simplifying Chemical Kinetics: Intrinsic Low-Dimensional Manifolds in Composition Space". In: *Combust. Flame* 88.3 (1992), pp. 239–264.

[131] Z. Ren and S. B. Pope. "Transport-Chemistry Coupling in the Reduced Description of Reactive Flows". In: *Combust. Theor. Model.* 11.5 (2007), pp. 715–739.

[132] H. Bongers, J. Van Oijen, and L. De Goey. "Intrinsic Low-Dimensional Manifold Method Extended with Diffusion". In: *Proc. Combust. Inst.* 29.1 (2002). Proceedings of the Combustion Institute, pp. 1371–1378.

[133] J. V. Oijen and L. D. Goey. "Modelling of Premixed Laminar Flames Using Flamelet-Generated Manifolds". In: *Combust. Sci. Technol.* 161.1 (2000), pp. 113–137.

[134] J. Oijen, van. "Flamelet-Generated Manifolds : Development and Application to Premixed Laminar Flames". English. PhD thesis. Technische Universiteit Eindhoven, 2002.

[135] E. V. Klapdor. "Simulation of Combustor-Turbine Interaction in a Jet Engine". English. PhD thesis. Darmstadt: Technische Universität Darmstadt, June 2011.

[136] N. Swaminathan and A. Parente. *Machine Learning and Its Application to Reacting Flows: ML and Combustion*. Springer Nature, 2023.

[137] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.

[138] J.-Y. Chen et al. "An Economical Strategy for Storage of Chemical Kinetics: Fitting in Situ Adaptive Tabulation with Artificial Neural Networks". In: *Proc. Combust. Inst.* 28.1 (2000), pp. 115–121.

[139] J. Blasco et al. "Modelling the Temporal Evolution of a Reduced Combustion Chemical System With an Artificial Neural Network". In: *Combust. Flame* 113.1 (1998), pp. 38–52.

[140] R. Laubscher. "Utilization of Artificial Neural Networks to Resolve Chemical Kinetics in Turbulent Fine Structures of an Advanced CFD Combustion Model". English. PhD thesis. Stellenbosch: Stellenbosch University, 2017.

[141] Z. Nikolaou, L. Vervisch, and P. Domingo. "Criteria to Switch From Tabulation to Neural Networks in Computational Combustion". In: *Combust. Flame* 246 (2022), p. 112425.

[142] F. Christo et al. “An Integrated PDF/neural Network Approach for Simulating Turbulent Reacting Systems”. In: *Symp. (Int.) Combust.* 26.1 (1996), pp. 43–48.

[143] B. A. Sen and S. Menon. “Linear Eddy Mixing Based Tabulation and Artificial Neural Networks for Large Eddy Simulations of Turbulent Flames”. In: *Combust. Flame* 157.1 (2010), pp. 62–74.

[144] W. Y. Peng and N. H. Pinkowski. *Efficient and Accurate Time-Integration of Combustion Chemical Kinetics Using Artificial Neural Networks*. 2017.

[145] R. Mao et al. “DeepFlame: A Deep Learning Empowered Open-Source Platform for Reacting Flow Simulations”. In: *Comput. Phys. Commun.* 291 (2023), p. 108842.

[146] F. A. Döppel and M. Votsmeier. “Efficient Neural Network Models of Chemical Kinetics Using a Latent Asinh Rate Transformation”. In: *React. Chem. Eng.* 8 (10 2023), pp. 2620–2631.

[147] T. Zhang et al. “A Multi-Scale Sampling Method for Accurate and Robust Deep Neural Network to Predict Combustion Chemical Kinetics”. In: *Combust. Flame* 245 (2022), p. 112319.

[148] Z.-Q. J. Xu et al. *Solving Multiscale Dynamical Systems by Deep Learning*. 2024. arXiv: `2401.01220 [math.NA]`.

[149] C. D. J Blasco N Fueyo and J.-Y. Chen. “A Self-Organizing-Map Approach to Chemistry Representation in Combustion Applications”. In: *Combust. Theor. Model.* 4.1 (2000), pp. 61–76.

[150] M. Haghshenas et al. “Acceleration of Chemical Kinetics Computation with the Learned Intelligent Tabulation (LIT) Method”. In: *Energies* 14.23 (2021).

[151] A. Almeldein and N. Van Dam. “Accelerating Chemical Kinetics Calculations With Physics Informed Neural Networks”. In: *ASME J. Eng. Gas Turbines Power* 145.9 (July 2023), p. 091008.

[152] S. Goswami et al. “Learning Stiff Chemical Kinetics Using Extended Deep Neural Operators”. In: *Comput. Methods Appl. Mech. Eng.* 419 (2024), p. 116674.

[153] L. Lu et al. “Learning Nonlinear Operators Via DeepONet Based on the Universal Approximation Theorem of Operators”. In: *Nat. Mach. Intell.* 3.3 (Mar. 2021), pp. 218–229.

[154] J. C. Sutherland and A. Parente. “Combustion Modeling Using Principal Component Analysis”. In: *Proc. Combust. Inst.* 32.1 (2009), pp. 1563–1570.

[155] H. Mirgolbabaei and T. Echekki. “Nonlinear Reduction of Combustion Composition Space with Kernel Principal Component Analysis”. In: *Combust. Flame* 161.1 (2014), pp. 118–126.

[156] E. Armstrong and J. C. Sutherland. “Reduced-Order Modeling with Reconstruction-Informed Projections”. In: *Combust. Flame* 259 (2024), p. 113119.

[157] M. R. Malik et al. “Combustion Modeling Using Principal Component Analysis: A Posteriori Validation on Sandia Flames D, E and F”. In: *Proc. Combust. Inst.* 38.2 (2021), pp. 2635–2643.

[158] K. Zdybał, A. Parente, and J. C. Sutherland. “Improving Reduced-Order Models Through Nonlinear Decoding of Projection-Dependent Outputs”. In: *Patterns* 4.11 (2023), p. 100859.

[159] V. Vijayarangan et al. "A Data-Driven Reduced-Order Model for Stiff Chemical Kinetics Using Dynamics-Informed Training". In: *Energy and AI* 15 (2024), p. 100325.

[160] A. Kumar and T. Echekki. "Combustion Chemistry Acceleration with DeepONets". In: *Fuel* 365 (2024), p. 131212.

[161] P. Zhang and R. Sankaran. "Autoencoder Neural Network for Chemically Reacting Systems". In: *J. Machine Learning for Modeling and Computing* 3.4 (2022), pp. 1–28.

[162] Grassi, T. et al. "Reducing the Complexity of Chemical Networks Via Interpretable Autoencoders". In: *A&A* 668 (2022), A139.

[163] K. S. Jung et al. "On the Application of Principal Component Transport for Compression Ignition of Lean Fuel/air Mixtures Under Engine Relevant Conditions". In: *Combust. Flame* 260 (2024), p. 113204.

[164] A. Kumar et al. "Acceleration of Turbulent Combustion DNS Via Principal Component Transport". In: *Combust. Flame* 255 (2023), p. 112903.

[165] L. He and M. L. G. Oldfield. "Unsteady Conjugate Heat Transfer Modeling". In: *ASME J. Turbomach.* 133.3 (Nov. 2010), p. 031022.

[166] L. He. "Closely Coupled Fluid-Solid Interface Method with Moving-Average for LES Based Conjugate Heat Transfer Solution". In: *Int. J. Heat Fluid Flow* 79 (2019), p. 108440.

[167] C. M. Bishop and H. Bishop. *Deep Learning: Foundations and Concepts*. Springer Nature, 2023.

[168] I. Sobol'. "On the Distribution of Points in a Cube and the Approximate Evaluation of Integrals". In: *USSR Comput. Math. & Math. Phys.* 7.4 (1967), pp. 86–112.

[169] X. Han et al. "An Improved Approach Towards More Robust Deep Learning Models for Chemical Kinetics". In: *Combust. Flame* 238 (2022), p. 111934.

[170] K. Zdybał et al. "Cost Function for Low-Dimensional Manifold Topology Assessment". In: *Sci. Rep.* 12.1 (Aug. 2022), p. 14496.

[171] F. Chollet et al. *Keras*. `https://keras.io`. 2015.

[172] Martín Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015.

[173] R. A. Saleh and A. K. M. E. Saleh. *Statistical Properties of the Log-Cosh Loss Function Used in Machine Learning*. 2024. arXiv: `2208.04564 [stat.ML]`.

[174] J. D. Anderson. *Hypersonic and High Temperature Gas Dynamics*. Aiaa, 1989.

[175] J. D. Denton. "The Calculation of Three-Dimensional Viscous Flow Through Multistage Turbomachines". In: *ASME J. Turbomach.* 114.1 (Jan. 1992), pp. 18–26.

[176] ANSYS. *ANSYS Fluent - CFD Software | ANSYS*. 2016. URL: `http://www.ansys.com/products/fluids/ansys-fluent`.

[177] G. Mudalige et al. "OP2: An Active Library Framework for Solving Unstructured Mesh-Based Applications on Multi-Core and Many-Core Architectures". In: *2012 Innovative Parallel Computing (InPar)*. 2012, pp. 1–12.

[178] D. Meyer et al. "Machine Learning Emulation of Urban Land Surface Processes". In: *J. Adv. Model. Earth Syst.* 14.3 (2022), e2021MS002744.

[179] P. Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nat. Methods* 17 (2020), pp. 261–272.

[180] K. Wan et al. "Chemistry Reduction Using Machine Learning Trained From Non-Premixed Micro-Mixing Modeling: Application to DNS of a Syngas Turbulent Oxy-Flame with Side-Wall Effects". In: *Combust. Flame* 220 (2020), pp. 119–129.

[181] R. T. Q. Chen et al. *Neural Ordinary Differential Equations*. 2019. arXiv: `1806.07366 [cs.LG]`.

[182] O. Owoyele and P. Pal. "ChemNODE: A Neural Ordinary Differential Equations Framework for Efficient Chemical Kinetic Solvers". In: *Energy and AI* 7 (2022), p. 100118.

[183] S. Bansude, F. Imani, and R. Sheikhi. "Performance Assessment of Chemical Kinetics Neural Ordinary Differential Equations in Pairwise Mixing Stirred Reactor". In: *ASME Open J. Engineering* 2 (Jan. 2023), p. 021008.

[184] S. Kim et al. "Stiff Neural Ordinary Differential Equations". In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 31.9 (Sept. 2021), p. 093122.

[185] C. Zhang et al. "Dynamic Coking Simulation of Supercritical N-Decane in Circular Tubes". In: *Fuel* 331 (2023), p. 125859.

[186] C. Scherding et al. "Data-Driven Framework for Input/output Lookup Tables Reduction: Application to Hypersonic Flows in Chemical Nonequilibrium". In: *Phys. Rev. Fluids* 8 (2 Feb. 2023), p. 023201.

[187] K. Peerenboom et al. "Dimension Reduction of Non-Equilibrium Plasma Kinetic Models Using Principal Component Analysis". In: *Plasma Sources Sci. Technol.* 24.2 (Jan. 2015), p. 025004.

[188] T. Berkemeier et al. "Accelerating Models for Multiphase Chemical Kinetics Through Machine Learning with Polynomial Chaos Expansion and Neural Networks". In: *Geosci. Model Dev.t* 16.7 (2023), pp. 2037–2054.

[189] J. L. A. M. Hendrix, A. J. Louca, and Y. Miguel. "Using a Neural Network Approach to Accelerate Disequilibrium Chemistry Calculations in Exoplanet Atmospheres". In: *Monthly Notices of the Royal Astronomical Society* 524.1 (June 2023), pp. 643–655.

[190] J. Zhang et al. "Combined Catalytic and Pyrolytic Coking Model for Steam Cracking of Hydrocarbons". In: *Ind. Eng. Chem. Res.* 61.11 (2022), pp. 3917–3927.

[191] D. Hendrycks and K. Gimpel. *Gaussian Error Linear Units (GELUs)*. 2023. arXiv: 1606.08415 [cs.LG].

[192] I. Loshchilov and F. Hutter. *Decoupled Weight Decay Regularization*. Oct. 2019. arXiv: 1410.5093 [cs.LG].