

Master's thesis and Internship : Modelling and Integration of a Hydrocracking Unit for Fischer-Tropsch Syncrude in a Power-to-Kerosene Process

Auteur : Laurent, Edwine

Promoteur(s) : Léonard, Grégoire

Faculté : Faculté des Sciences appliquées

Diplôme : Master : ingénieur civil en génie de l'énergie à finalité spécialisée en Energy Conversion

Année académique : 2025-2026

URI/URL : <http://hdl.handle.net/2268.2/26189>

Avertissement à l'attention des usagers :

Tous les documents placés en accès ouvert sur le site le site MatheO sont protégés par le droit d'auteur. Conformément aux principes énoncés par la "Budapest Open Access Initiative"(BOAI, 2002), l'utilisateur du site peut lire, télécharger, copier, transmettre, imprimer, chercher ou faire un lien vers le texte intégral de ces documents, les disséquer pour les indexer, s'en servir de données pour un logiciel, ou s'en servir à toute autre fin légale (ou prévue par la réglementation relative au droit d'auteur). Toute utilisation du document à des fins commerciales est strictement interdite.

Par ailleurs, l'utilisateur s'engage à respecter les droits moraux de l'auteur, principalement le droit à l'intégrité de l'oeuvre et le droit de paternité et ce dans toute utilisation que l'utilisateur entreprend. Ainsi, à titre d'exemple, lorsqu'il reproduira un document par extrait ou dans son intégralité, l'utilisateur citera de manière complète les sources telles que mentionnées ci-dessus. Toute utilisation non explicitement autorisée ci-avant (telle que par exemple, la modification du document ou son résumé) nécessite l'autorisation préalable et expresse des auteurs ou de leurs ayants droit.

Modelling and Integration of a Hydrocracking Unit for Fischer-Tropsch Syncrude in a Power-to-Kerosene Process

LAURENT Edwine

Thesis presented to obtain the degree of :
Master of Science in Energy Engineering

Thesis supervisor :
LEONARD Grégoire

Academic year: **2025 - 2026**

Acknowledgments

This master's thesis marks the end of five years of study that have been as rich in learning as in friendships, and its completion would not have been possible without the help and support of many people.

I would first like to thank the professors of the Faculty of Applied Sciences for the knowledge they have transmitted over these years, which provided the foundation on which this work is built. A particular thanks goes to Professor G. Léonard for supervising this thesis and for the critical perspective he brought to each of our exchanges throughout the project.

I am especially grateful to doctoral researcher A. Rouxhet who, despite being in the final year of his own thesis, consistently made time to follow my reasoning, to read and comment on this work before its submission. His involvement, our discussions, and his remarks have really improved this work.

I would also like to thank the other doctoral students of Professor Léonard's team for their warm welcome in their office. I am grateful to Professor P. Dewallef, Professor D. Toye and V. Sang Sefidi for agreeing to be part of my jury and for the time they devoted to reading this report. I also thank Professor D. Toye for the discussions and for answering my questions regarding the sizing of the reactor.

Finally, I would like to thank my family and my friends for their encouragement throughout the moments of doubt and intense work that any thesis brings. A particular thought goes to my parents, who gave me every opportunity and every key that made all of this possible, and to my godfather, who first passed on to me a taste for engineering. I owe a special thanks to my boyfriend, for his daily presence, his patience, and the countless conversations that helped me find my way through the most demanding moments of this work.

To all of them, I am sincerely grateful.

Abstract

Aviation is among the hardest sectors to decarbonise: long-haul flight depends on high-energy-density liquid fuels for which direct electrification offers no realistic substitute. The Power-to-Liquids route, in which captured CO₂ is combined with electrolytic hydrogen to synthesise drop-in jet fuel, represents a relevant pathway for reducing aviation's dependence on fossil kerosene. In the low-temperature Fischer-Tropsch route over a cobalt-based catalyst, syngas is converted into a broad hydrocarbon distribution that is, however, weighted towards the heavy C₁₇₊ paraffinic wax fraction rather than the C₈-C₁₆ kerosene cut targeted by the process. Converting this heavy fraction into kerosene-range products therefore requires a dedicated upgrading step, and hydrocracking over a bifunctional catalyst is the technology retained for this purpose. The present work develops a hydrocracking reactor model and integrates it downstream of the Fischer-Tropsch (FT) unit of the Power-to-Kerosene process under development at the University of Liège (ULiège) within the Neutral-Kero-Lime project, completing the upstream chain previously modelled by A. Rouxhet and A. Morales [1,2].

Three modelling approaches were evaluated successively. A stoichiometric model and a thermodynamic-equilibrium model were both shown to be structurally unable to reproduce the experimental hydrocracking behaviour, confirming that the reaction is kinetically rather than thermodynamically controlled. The Langmuir-Hinshelwood-Hougen-Watson kinetic model of Pellegrini et al. (2007) [3] was consequently retained, because it resolves the product stream at the component level, a resolution required to track the C₈-C₁₆ cut without committing to a fixed lump definition, under the simplifying assumption of a single vapour phase combined with midpoint cracking. Two pieces of information omitted from the original publication had to be recovered in Python before the model could be applied: the inlet distribution, identified through hypothesis testing as corresponding to the distribution later reported by Pellegrini et al. (2008) [4], and the final optimised kinetic parameters, reconstructed by regularised parameter estimation against the published model curves.

The validated model was then applied to the ULiège FT hydrocarbon stream after olefin-hydrogenation pre-treatment, with the objective of identifying the temperature, pressure, weight hourly space velocity (WHSV) and hydrogen-to-wax ratio that maximise the C₈-C₁₆ mass fraction at the reactor outlet. Within the Pellegrini calibration domain, the procedure converges at $T = 369.3^\circ\text{C}$, $P = 60\text{ bar}$, $\text{WHSV} = 3.00\text{ kg}_{\text{n-C}_{10}}\text{-in kg}_{\text{cat}}^{-1}\text{ h}^{-1}$ and $\text{H}_2/\text{wax} = 0.150\text{ kg/kg}$, yielding 52.0 wt-% of C₈-C₁₆ at the outlet for a C₁₇₊ mass conversion of 55.9 %. At the pilot-scale throughput of the ULiège installation, applying standard plug-flow design criteria to this operating point yields a preliminary trickle-bed reactor geometry of approximately 26 mm internal diameter and 132 mm bed height, holding close to 38 g of bifunctional Pt/SiO₂-Al₂O₃ catalyst. Three of the four coordinates lie on the upper bound of the calibration window, indicating that the model response would continue to climb beyond the calibrated range. Dimensional sensitivity analyses around this point expose two limitations of the model: the H₂/wax ratio acts on the model exclusively through the hydrogen partial pressure, so the genuine two-phase character of the system is not captured; and the midpoint cracking stoichiometry underestimates the iso-paraffin yield relative to the experimental data. A comparison of the recommended conditions across feed selection (full C₁-C₇₀₊ stream versus C₁₇₊ fraction only), exploration domain (author calibration versus extended literature range) and optimality criterion (strict versus relaxed economic ranking) further isolates the contribution of each modelling choice. One caveat conditions the reading: the kinetic parameters, calibrated on the Pellegrini C₄-C₇₀ feedstock and not refitted to the ULiège stream (for which no experimental data exist), make the reported yields and conditions order-of-magnitude indications and a transferable methodology, not definitive design values, though a check on the calibration feed itself suggests the trends carry over between feeds.

Several perspectives follow naturally. At the reactor scale, replacing the single-phase partial-pressure assumption with a proper vapour-liquid description and the midpoint cracking with a breakage-probability stoichiometry would address the two structural limitations identified by the sensitivity analyses. At the process scale, integrating the model into the full Power-to-Kerosene flowsheet in Aspen Plus would enable a techno-economic assessment of the recommended operating point.

Contents

Introduction	1
Objective	1
1 State of the art	2
1.1 The Power-to-X concept	2
1.1.1 Power-to-Gas	2
1.1.1.1 Power-to-Hydrogen	2
1.1.1.2 Power-to-Methane	3
1.1.2 Power-to-Liquids	3
1.1.2.1 Power-to-Methanol	4
1.1.2.2 PtL production via the Fischer-Tropsch pathway	4
1.1.2.2.1 The reverse water-gas shift reaction	5
1.1.2.2.2 Fischer-Tropsch synthesis	5
1.1.2.2.3 Upgrading	6
1.1.2.2.4 Example: the ULiège process	6
1.1.3 Power-to-Chemicals	6
1.1.4 Power-to-Heat	7
1.2 Aviation kerosene: production routes and fuel specifications	7
1.2.1 Conventional kerosene production	7
1.2.2 Sustainable aviation fuels	8
1.2.2.1 Quality specifications and standards	8
1.3 Upgrading of Fischer-Tropsch syncrude	9
1.3.1 Upgrading technology selection	9
1.3.2 Justification for the use of hydrocracking	10
1.3.3 Possible cracking routes	10
1.3.3.1 Thermal cracking	10
1.3.3.2 Fluid catalytic cracking (FCC)	11
1.3.4 Comparison with conventional hydrocracking	11
1.4 Hydrocracking of LTFT syncrude	12
1.4.1 Reaction mechanism and catalysts	12
1.4.1.1 Bifunctional reaction pathway	13
1.4.1.2 Ideal and non-ideal bifunctional catalysts	13
1.4.1.3 Catalyst materials	14
1.4.1.3.1 Acidic supports	14
1.4.1.3.2 Metallic components	14
1.4.1.3.3 Resistance to deactivation	15
1.4.2 Stoichiometry	15
1.4.2.1 Hydroisomerization stoichiometry	15
1.4.2.2 Cracking step stoichiometry	16

1.4.2.2.1	Midpoint approximation	17
1.4.2.2.2	Breakage probability function	17
1.4.3	Kinetics	18
1.4.3.1	Overview of kinetic modelling approaches	18
1.4.3.2	The Pellegrini-Gambaro research series: common framework	20
1.4.3.3	Article-by-article development	20
1.4.3.3.1	Pellegrini et al. (2004): The lumped model [5]	20
1.4.3.3.2	Pellegrini et al. (2007): The all-components model [3]	21
1.4.3.3.3	Pellegrini et al. (2008): Introduction of vapour-liquid equilibrium [4]	21
1.4.3.3.4	Gamba et al. (2009): Introduction of a breakage probability function [6]	21
1.4.3.3.5	Gambaro et al. (2011): Complete LHHW rate expression [7]	22
1.4.4	Operating conditions	22
1.4.5	Reactor configurations	23
1.4.5.1	Trickle-bed as the industrial standard	23
1.4.5.2	Operating principle	23
1.4.5.3	Suitability for the Co-LTFT wax feed	24
2	Hydrocracking reactor modelling	25
2.1	Feed stream characterisation	26
2.1.1	Origin and composition of the feed	26
2.1.2	Lumping strategy	27
2.1.3	Olefin hydrogenation pre-treatment	27
2.2	Model selection and implementation framework	28
2.2.1	Choice of Python as modelling tool	28
2.2.2	Selection of the all-components model of Pellegrini et al. (2007)	29
2.3	Kinetic model implementation	30
2.3.1	Material balance equations	30
2.3.2	Kinetic parameter estimation	34
2.3.3	Thermodynamic assumption	36
2.4	Operating conditions and optimal set identification	37
2.4.1	Operating window and explored cases	37
2.4.2	Search strategy for the kerosene-maximising operating conditions	39
2.5	Reactor sizing model for LTFT wax hydrocracking	43
2.5.1	Catalyst mass from WHSV	43
2.5.2	Packed-bed volume	43
2.5.3	Reactor geometry	44
2.5.4	Hydrogen co-feed	45
2.5.5	Heat management	45
3	Results and discussions	46
3.1	Kinetic model validation	47
3.1.1	Identification of the Pellegrini (2007) inlet distribution	47
3.1.2	Stoichiometric model	50
3.1.3	Thermodynamic model	52
3.1.4	LHHW kinetic model: parameter estimation and model reproduction	54
3.2	Application to the ULiège inlet stream	59
3.2.1	Inlet composition after olefin hydrogenation	59
3.2.2	Kinetic profiles with ULiège inlet	61
3.2.2.1	Optimal operating conditions	61
3.2.2.2	Kinetic plots	62
3.2.2.3	Sensitivity analysis	65

TABLE OF CONTENTS

3.2.2.3.1	Two-dimensional sensitivity slices	65
3.2.2.3.2	Reactor profiles under varying operating conditions	67
3.2.2.3.3	One-dimensional sweeps	69
3.2.2.4	Thermodynamic admissibility under perturbations	71
3.2.3	Reactor sizing	72
3.2.3.1	Catalyst mass and hydrogen co-feed	72
3.2.3.2	Bulk density of the catalytic bed	73
3.2.3.3	Packed-bed volume and reactor geometry.	73
3.2.3.4	Synthesis and interpretation	74
3.3	Optimal operating conditions with other considerations	74
3.3.1	Strict-optimality scenario	75
3.3.2	Relaxed-optimality scenario	76
3.3.3	Synthesis	77
3.3.4	Sensitivity of the optimum to the feed composition	78
4	Conclusions and perspectives	80
4.1	Conclusions	80
4.2	Perspectives	81
	Appendices	83
	Disclaimer on AI	87
	Bibliography	93

List of Figures

1.1	Power-to-Kerosene process following the Fischer-Tropsch route [8].	5
1.2	The conventional route for jet fuel production from crude oil [9].	7
1.3	Main production routes for sustainable aviation fuel [10].	8
1.4	Technology selection for the refining of syncrude to jet fuel [11].	10
1.5	Assumed product distribution for the hydrocracking of $C_{60}H_{122}$	18
1.6	Schematic representation of a trickle-bed reactor operated in co-current gas-liquid down-flow mode.	24
2.1	Simplified block flow diagram of the relevant section of the ULiège process and the hydrocracking unit modelled in this work.	26
2.2	Initial molar-fraction distribution of the ULiège hydrocarbon inlet stream.	27
2.3	Final distribution obtained by Pellegrini et al. under conditions $T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $H_2/\text{wax} = 0.1275\text{ kg/kg}$, $\text{WHSV} = 2.5\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$ using the final kinetic constants expressed as functions of carbon number [3].	35
2.4	Recovery ratio in vapour phase for each component, plotted by Pellegrini et al. [3].	36
2.5	Four-phase search procedure used to identify the recommended operating point.	42
3.1	Inlet distribution constructed from the lumped feedstock composition reported by Pellegrini et al. (2004) [5], assuming a uniform mass distribution of species within each carbon-number cut.	47
3.2	Outlet carbon-number distributions simulated with the lumped 2004 inlet compared with the lumped-model curves published by Pellegrini et al. (2007) [3], under conditions ($T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $H_2/\text{wax} = 0.1275\text{ kg/kg}$, $\text{WHSV} = 2.5\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$).	48
3.3	Parity plots of the outlet molar fractions simulated with the component-level 2004 inlet against the lumped-model curves of Pellegrini et al. (2007) [3,5].	48
3.4	n-Paraffin inlet distribution reported by Pellegrini et al. (2008) [4].	49
3.5	Outlet carbon-number distributions simulated with the component-level 2008 inlet compared with the lumped-model curves published by Pellegrini et al. (2007) [3], under conditions ($T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $H_2/\text{wax} = 0.1275\text{ kg/kg}$, $\text{WHSV} = 2.5\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$).	50
3.6	Parity plots of the outlet molar fractions simulated with the component-level 2008 inlet against the lumped-model curves of Pellegrini et al. (2007) [3,4].	50
3.7	Comparison between the simulated outlet carbon-number distribution and the experimental data of Pellegrini et al. (2007) [3] after one stoichiometric cracking pass with complete conversion of all C_{5+} species.	51
3.8	Comparison between the simulated outlet carbon-number distribution and the experimental data of Pellegrini et al. (2007) [3] after two successive stoichiometric cracking passes with complete conversion at each pass.	52
3.9	Outlet carbon-number distribution predicted by the RGibbs model (logarithmic scale) compared with the initial distribution and the experimental outlet of Pellegrini et al. (2007) [3], at $T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $H_2/\text{wax} = 0.1275\text{ kg/kg}$	53

3.10	Outlet carbon-number distribution predicted by the REquil model with midpoint cracking stoichiometry (logarithmic scale) compared with the initial distribution and the experimental outlet of Pellegrini et al. (2007) [3], at $T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $\text{H}_2/\text{wax} = 0.1275\text{ kg/kg}$	53
3.11	Outlet carbon-number distribution predicted by the REquil model with the breakage probability cracking stoichiometry (logarithmic scale) compared with the initial distribution and the experimental outlet of Pellegrini et al. (2007) [3], at $T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $\text{H}_2/\text{wax} = 0.1275\text{ kg/kg}$	54
3.12	Outlet molar fraction distributions of n-paraffins (left) and iso-paraffins (right) simulated before and after parameter optimisation against the model curves of Pellegrini et al. (2007) [3], under conditions ($T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $\text{H}_2/\text{wax} = 0.1275\text{ kg/kg}$, $\text{WHSV} = 2.5\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$).	55
3.13	Carbon-number dependence of the kinetic parameters before optimisation (initial expressions of Pellegrini et al. (2007)) and after optimisation targeting the Pellegrini model curves.	56
3.14	Outlet molar fraction distributions of n-paraffins (left) and iso-paraffins (right) obtained after parameter optimisation against the experimental data of Pellegrini et al. (2007) [3], under conditions ($T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $\text{H}_2/\text{wax} = 0.1275\text{ kg/kg}$, $\text{WHSV} = 2.5\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$).	57
3.15	Carbon-number dependence of the kinetic parameters before optimisation (initial expressions of Pellegrini et al. (2007)) and after optimisation targeting the experimental data.	58
3.16	Molar composition of the FT hydrocarbon product stream of ULiège before and after olefin hydrogenation.	60
3.17	Comparison of the ULiège hydrocracking reactor inlet stream after olefin hydrogenation with the $\text{C}_4\text{-C}_{70}$ n-paraffin inlet distribution of Pellegrini et al. (2007) [4].	61
3.18	Inlet and outlet carbon-number distributions at the optimal operating point ($T = 369.3\text{ }^{\circ}\text{C}$, $P = 60\text{ bar}$, $\text{WHSV} = 3.00\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$, $\text{H}_2/\text{wax} = 0.150$).	62
3.19	Reactor profiles as a function of the space time $\tau = 1/\text{WHSV}$ at the recommended operating point.	63
3.20	Outlet molar distribution of total paraffins (n + iso combined), compared with the feed.	64
3.21	Outlet molar distributions of n- and iso-paraffins at the recommended operating point.	64
3.22	Two-dimensional sensitivity slices of the $\text{C}_8 - \text{C}_{16}$ outlet mass fraction around the recommended operating point ($T = 369.3\text{ }^{\circ}\text{C}$, $P = 60\text{ bar}$, $\text{WHSV} = 3.00\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$, $\text{H}_2/\text{wax} = 0.150\text{ kg/kg}$).	65
3.23	Mass conversion of the C_{17+} fraction as a function of the space time $\tau = 1/\text{WHSV}$, for variations of each operating variable around the recommended point. In every panel, the three variables not being swept are held at $T = 369.3\text{ }^{\circ}\text{C}$, $P = 60\text{ bar}$, $\text{WHSV} = 3.00\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$, $\text{H}_2/\text{wax} = 0.150\text{ kg/kg}$	68
3.24	Sweep of the temperature, $P = 60\text{ bar}$, $\text{WHSV} = 3.00\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$, $\text{H}_2/\text{wax} = 0.150\text{ kg/kg}$	69
3.25	Sweep of the total pressure, $T = 369.3\text{ }^{\circ}\text{C}$, $\text{WHSV} = 3.00\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$, $\text{H}_2/\text{wax} = 0.150\text{ kg/kg}$	70
3.26	Sweep of the WHSV, $T = 369.3\text{ }^{\circ}\text{C}$, $P = 60\text{ bar}$, $\text{H}_2/\text{wax} = 0.150\text{ kg/kg}$	71
3.27	Sweep of the hydrogen-to-wax mass ratio, $T = 369.3\text{ }^{\circ}\text{C}$, $P = 60\text{ bar}$, $\text{WHSV} = 3.00\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$	71
3.28	Inlet overall vapour mole fraction (VFRAC) and liquid mole fraction (LFRAC) computed in Aspen Plus with the SRK equation of state, at $\text{H}_2/\text{wax} = 0.150$	72
3.29	Schematic representation of the cylindrical packed-bed reactor at the recommended operating point (full $\text{C}_1\text{-C}_{70+}$ ULiège feed, $T = 369.3\text{ }^{\circ}\text{C}$, $P = 60\text{ bar}$, $\text{WHSV} = 3.00\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$, $\text{H}_2/\text{wax} = 0.15\text{ kg/kg}$).	74

List of Tables

1.1	Carbon number distribution and typical compound classes of high-temperature Fischer-Tropsch and low-temperature Fischer-Tropsch products, excluding C_1 - C_2 hydrocarbons [12].	5
1.2	Selected ASTM property requirements for Jet A/A-1 and representative SPK-type SAF pathways [10].	9
1.3	Typical conditions for conventional and LTFT wax hydrocracking [11,13–16].	12
1.4	Range of operating conditions for Co-LTFT wax hydrocracking reported in the literature [5, 17].	22
2.1	Stream properties at the flash separator outlet before and after removal of non-reactive species.	27
2.2	Definition of variables and parameters appearing in the material balance equations.	33
2.3	Operating domains explored in this work.	37
3.1	Refined optima retained under the strict-optimality scenario.	75
3.2	Refined optima retained under the relaxed-optimality scenario.	76
3.3	Recommended operating conditions for the eight cases explored.	77
3.4	Recommended operating conditions for the ULiège and Pellegrini feeds, strict-precision scenario over the Pellegrini calibration domain.	79
A.1	Typical syncrude compositions for Fe-HTFT, Fe-LTFT and Co-LTFT Fischer-Tropsch syntheses [11].	83
A.2	Thermodynamic state (A.2a) and molar composition of the stream (A.2b) at the flash separator outlet, as provided by A. Rouxhet.	84
A.3	Kinetic parameter values used for the first model in Pellegrini et al. (2007) [3].	84
A.4	Molar fractions in the FT hydrocarbon stream after removal of C_1 - C_{16}	85
A.5	Initial molar-fraction distribution x_i of n-paraffins C_nH_{2n+2} from Pellegrini et al. (2008) [4].	85
A.6	Optimised parameters used to reproduce the Pellegrini et al. (2007) model.	86
A.7	Detailed molar fractions of Fischer-Tropsch products provided by A. Rouxhet.	86

Introduction

What if travelling without polluting became a reality? Today the idea is hard to believe when a Dublin-Bergamo flight that emits up to 100 kg of CO₂ per passenger can be booked for €19.99, the price of a pizza [18]. At that price, the temptation to fly is hard to resist.

Behind such fares lies a sector that has become indispensable to global mobility and trade, and one that is, at the same time, among the hardest to decarbonise. Hard to electrify knowing that replacing the energy stored in one kilogram of jet fuel currently requires roughly fifty kilograms of battery [10, 19]. Battery-electric and hydrogen propulsion may eventually serve short-range aviation, but long-haul flight, which dominates the sector's fuel consumption, remains tied to a high-energy-density liquid fuel. The relevant question is therefore not whether aircraft will keep burning liquid hydrocarbons, but whether those hydrocarbons can be produced without fossil carbon.

Aviation's climate footprint is both large and growing. In 2023, aviation was responsible for about 2.5% of global energy-related CO₂ emissions, close to 950 Mt, and the sector expanded faster than rail, road or shipping between 2000 and 2019 [19]. Each litre of fossil jet fuel releases roughly 2.5 kg of CO₂ on combustion [10], and once non-CO₂ effects such as contrails and nitrogen oxides are accounted for, the contribution of aviation to climate forcing is estimated to rise to 4-5% [20]. Regulators have responded accordingly: the International Civil Aviation Organization has adopted a net-zero target for international aviation by 2050, while the European Union's ReFuelEU Aviation regulation now imposes a minimum sustainable aviation fuel share, rising from 2% in 2025 to 6% in 2030, together with a dedicated sub-target for synthetic fuels [21].

Beyond the climate imperative, the current geopolitical context provides a second, independent argument for synthetic aviation fuels. In Europe, conventional kerosene is refined from imported crude oil, a geographically concentrated resource that leaves aviation exposed to supply disruptions and price volatility, whereas e-kerosene can be produced from renewable electricity, water and captured CO₂, inputs that can be sourced locally.

These synthetic fuels are produced by Power-to-Liquids processes, which convert renewable electricity into value-added liquid products, effectively storing it in chemical form. Water is electrolysed to hydrogen, captured CO₂ is reduced to carbon monoxide through the reverse water-gas shift reaction, and the resulting synthesis gas is converted into liquid hydrocarbons over a Fischer-Tropsch (FT) catalyst; when the targeted product is jet fuel, the route is referred to as Power-to-Kerosene. Because its carbon is recycled from captured CO₂ rather than extracted from crude oil, the fuel does not rely on fossil carbon, and FT-derived kerosene additionally burns more cleanly than crude-derived Jet A-1, producing markedly lower particulate emissions [11].

Fischer-Tropsch synthesis, however, does not yield a finished jet fuel on its own: it produces a broad mixture of hydrocarbons, weighted toward heavy paraffinic waxes that lie well above the C₈-C₁₆ kerosene range. Converting this heavy fraction into kerosene-range hydrocarbons requires a dedicated upgrading step, for which hydrocracking over a bifunctional catalyst is the technology retained.

The modelling of this hydrocracking step, and its integration into the Power-to-Kerosene process developed at the University of Liège, is the subject of the present work; its objectives are set out in the section that follows.

Objectives

As stated in the Introduction, aviation being a hard-to-abate sector, Power-to-Liquids appears as the solution to preserve the high energy density of kerosene while remaining carbon-neutral. It is within this perspective that Professor G. Léonard's team contributes to the Neutral-Kero-Lime project, dedicated to setting up a pilot installation for e-kerosene production at the University of Liège [1]. The project aims at designing the process from scratch: captured CO_2 and electrolytic hydrogen are first converted to CO through a reverse water-gas shift reactor, and the resulting synthesis gas is then fed to a Fischer-Tropsch reactor that yields a broad spectrum of hydrocarbon products, referred to as syncrude.

The upstream sections of this process, composed of the electrolyser, the reverse water-gas shift (RWGS) reactor and a cobalt-based low-temperature Fischer-Tropsch (LTFT) reactor, are already modelled in Aspen Plus and Aspen Custom Modeler, the RWGS and FT reactors having been developed by A. Rouxhet and A. Morales [2]. Its outlet stream, however, is weighted predominantly towards the heavy C_{17+} fraction, which lies outside the $\text{C}_8\text{-C}_{16}$ kerosene cut targeted by the process. A dedicated downstream hydrocracking unit is therefore required to upgrade this heavy fraction, and no model of it has yet been developed within the ULiège chain.

The objective pursued in this work is therefore to design and simulate a hydrocracking unit for the Fischer-Tropsch syncrude in Python. This includes the selection of a catalyst, the identification and implementation of an appropriate reaction stoichiometry and kinetics, the determination of the operating conditions (temperature, pressure, weight hourly space velocity and hydrogen-to-wax ratio) that maximise the $\text{C}_8\text{-C}_{16}$ mass fraction at the reactor outlet, and the derivation of an indicative reactor geometry from those conditions. Sensitivity analyses on these key operating variables complete the study, and particular attention is also given to the influence of the syncrude composition at the reactor inlet.

The thesis is organised in four chapters. Chapter 1 reviews the Power-to-X framework, the production routes and specifications of aviation kerosene, and the reaction mechanism, stoichiometry, kinetics, and reactor configurations relevant to LTFT wax hydrocracking. Chapter 2 describes the construction of the reactor model: feed characterisation, model selection, the component-level kinetic implementation, the operating-condition search procedure, and the reactor sizing methodology. Chapter 3 reports the results, proceeding from model validation to application to the ULiège stream and to the broader analysis of the operating space. Chapter 4 draws the conclusions and outlines perspectives for future work.

Chapter 1

State of the art

1.1 The Power-to-X concept

Nowadays, the primary environmental objective is to mitigate global warming by significantly reducing or achieving net-zero CO₂ emissions. A major trend in achieving this goal is the electrification of various sectors to leverage the growing abundance of renewable electricity. However, complete electrification is unfeasible for certain hard-to-abate sectors, and the integration of intermittent renewables necessitates enhanced grid stabilisation strategies.

The Power-to-X (PtX) concept addresses these challenges by converting excess electricity generated by renewable energy sources into value-added products X, such as gases, liquids, chemicals, or heat. It enables the production of alternative energy carriers for sectors in which direct electrification remains impractical, particularly aviation and heavy transport, where high-energy-density liquid and gaseous fuels are still essential. Moreover, PtX facilitates the large-scale integration of renewable electricity by providing a means for long-term energy storage and grid balancing [22].


1.1.1 Power-to-Gas

Power-to-Gas (PtG) refers to the conversion of surplus electrical energy into gaseous energy carriers, enabling the large-scale integration of variable renewable electricity into the existing gas infrastructure. Two main conversion pathways can be distinguished among others: Power-to-Hydrogen (PtH), which produces hydrogen via water electrolysis, and Power-to-Methane (PtM), which further converts hydrogen and CO₂ into synthetic natural gas (SNG) through a methanation reaction.

1.1.1.1 Power-to-Hydrogen

Hydrogen is a versatile energy carrier with a high gravimetric energy density of approximately 120 MJ/kg_{H₂}, which is nearly three times that of conventional hydrocarbon fuels [23]. When used in fuel cells or combustion systems, it produces no direct CO₂ emissions, making it an attractive candidate for decarbonising a wide range of sectors. However, the environmental benefit of hydrogen depends strongly on how it is produced. Today, hydrogen production is overwhelmingly based on fossil fuels: global consumption reached 95–100 Mt in 2024, of which less than 0.1% was produced from low-carbon sources. As a result, hydrogen production is currently responsible for approximately 2.5% of global CO₂ emissions, equivalent to around 1 Gt per year [24]. This stark contrast between hydrogen's decarbonisation potential and its actual carbon footprint underscores the urgency of transitioning toward low-carbon production pathways.

Power-to-Hydrogen specifically refers to the production of hydrogen via water electrolysis, described by Eq. (1.1), powered by renewable electricity [25].




The three main commercially available electrolyser technologies are alkaline electrolyzers (AE), proton exchange membrane electrolyzers (PEM), and solid oxide electrolyser cells (SOEC) [26, 27]. Alkaline electrolyzers are the most mature and cost-effective option for large-scale, steady-state hydrogen production, having been deployed industrially for decades. PEM electrolyzers offer higher current densities, a compact design, and a fast dynamic response, making them particularly well-suited for coupling with variable renewable energy sources. SOEC systems operate at elevated temperatures (typically 700–900°C) and can achieve higher theoretical conversion efficiencies. Notably, SOECs can also perform co-electrolysis of H₂ and CO₂ to produce syngas directly, which opens a particularly relevant pathway toward synthetic fuel production. However, SOEC technology remains less commercially mature and faces durability challenges that limit its large-scale deployment [26, 28].

Despite its potential, green hydrogen faces significant economic barriers. Current production costs range from approximately 1.5 to 6 times those of conventional fossil-based hydrogen, depending on local electricity prices, electrolyser capital costs, capacity factors, and operating conditions [29]. Ongoing reductions in renewable electricity costs and electrolyser manufacturing are expected to improve competitiveness, but green hydrogen remains expensive at present. Beyond direct use in fuel cells or industrial processes, hydrogen serves as a key intermediate in the broader Power-to-X framework. Its combination with CO or CO₂ enables the synthesis of syngas, methane, methanol, and synthetic liquid fuels. For aviation in particular, the direct use of liquid hydrogen remains constrained by its very low volumetric energy density (8.5 MJ/L_{H₂} for the liquefied H₂ versus 35 MJ/L for conventional jet fuel [11, 23]), which would require a complete redesign of aircraft fuel systems and airport infrastructure [10]. This limitation motivates the conversion of green hydrogen into energy-dense liquid hydrocarbons via Power-to-Liquids pathways, as developed in Section 1.1.2.

1.1.1.2 Power-to-Methane

Power-to-Methane (PtM) extends the Power-to-Hydrogen pathway by converting the produced hydrogen and a CO₂ source into synthetic natural gas through the catalytic methanation reaction, described by Eq. (1.2), also known as the Sabatier reaction [30].



The reaction is exothermic and thermodynamically favoured at low temperatures and high pressures [30]. The primary advantage of PtM over direct hydrogen utilisation lies in its compatibility with the existing natural gas infrastructure. Unlike hydrogen, which requires dedicated pipelines and storage facilities, SNG can be injected directly into the natural gas grid without modification, enabling large-scale, long-term energy storage [26]. Several demonstration projects have validated this approach, including the Audi e-gas plant in Germany, which has been producing SNG from wind-powered electrolysis and biogenic CO₂ since 2013 [30]. However, the PtM process introduces an additional conversion step compared to PtH, resulting in energy losses that reduce the efficiency of the system. Furthermore, the combustion of methane, even of synthetic origin, still releases CO₂, which limits its contribution to deep decarbonisation unless the CO₂ feedstock is itself of non-fossil origin (e.g. captured from the atmosphere or from biogenic sources) [30]. These limitations, combined with the inherently low volumetric energy density of gaseous energy carriers, motivate the development of Power-to-Liquid pathways, as discussed in the following section.

1.1.2 Power-to-Liquids

Power-to-Liquids (PtL) refers to the production of synthetic liquid fuels from renewable electricity through an integrated process chain encompassing water electrolysis, carbon capture, synthesis gas generation, and catalytic fuel synthesis with subsequent product upgrading [28].

Compared to gaseous energy carriers such as hydrogen or synthetic methane, liquid fuels offer substantially higher volumetric energy densities, which makes them particularly well-suited for applications where weight and volume constraints are critical, most notably in aviation and heavy transport. Beyond

energy density, PtL-derived fuels present two decisive advantages for large-scale deployment. First, they are compatible with existing fuel distribution infrastructure, aircraft engines, and refuelling systems, thereby avoiding the prohibitive capital costs associated with a complete infrastructure overhaul.

Second, when produced from renewable electricity and atmospheric CO₂, PtL fuels can in principle achieve carbon neutrality over their life cycle, since the CO₂ released by combustion is balanced by the CO₂ originally withdrawn from the atmosphere [10].

While PtL encompasses a broad range of products, including for example dimethyl ether (DME), alcohols, and synthetic paraffinic kerosene, the following focuses on the two pathways most relevant to SAF production: the methanol-to-jet (MTJ) pathway, which proceeds via methanol as an intermediate, and the Fischer–Tropsch pathway, which converts syngas directly into a broad spectrum of hydrocarbons. Both are examined in the following sections.

1.1.2.1 Power-to-Methanol

Methanol is a versatile liquid energy carrier and a key intermediate in the broader Power-to-Liquids framework. In a Power-to-Methanol scheme, methanol is synthesised catalytically from CO₂ and renewable H₂ over Cu/ZnO/Al₂O₃ catalysts, typically operating at temperatures of 200–300 °C and pressures of 50–100 bar [10]. The overall conversion proceeds through three coupled reactions (1.3), (1.4) and (1.5).



A notable advantage of this pathway in the PtL context is that methanol can be synthesised directly from CO₂, without requiring a prior reverse water-gas shift stage or co-electrolysis step to generate CO. Each of the individual subprocesses of the subsequent methanol-to-jet (MTJ) chain, namely methanol-to-olefins conversion, oligomerization, and hydrogenation, is already operational at industrial scale in various plants and refineries. However, no commercial-scale implementation of a fully integrated MTJ process exists to date, reflecting its relatively low technology readiness level (TRL) [10, 28].

1.1.2.2 PtL production via the Fischer–Tropsch pathway

The Fischer–Tropsch pathway is currently the most mature and best-validated route for the production of synthetic aviation fuel from renewable electricity. Originally developed in Germany for producing liquid fuels from coal, the process has since been adapted to operate from a broad range of carbon feedstocks, including CO₂ captured from the atmosphere or from industrial flue gases [11]. In a Power-to-Liquids context, the process chain consists of three main stages, illustrated in Figure 1.1: the conversion of CO₂ and renewable H₂ into synthesis gas via the reverse water-gas shift reaction, the catalytic polymerisation of syngas into a hydrocarbon syncrude via the Fischer–Tropsch synthesis, and the upgrading of this syncrude into jet fuel fractions complying with aviation specifications. A particular focus is placed on this pathway in the present work, as the Fischer–Tropsch syncrude constitutes the feed stream to the hydrocracking reactor studied subsequently.





Figure 1.1: Power-to-Kerosene process following the Fischer-Tropsch route [8].

1.1.2.2.1 The reverse water-gas shift reaction

The first stage of the FT process chain consists of converting (captured) CO_2 and (renewable) H_2 into carbon monoxide via the reverse water-gas shift (RWGS) reaction (1.6).



The reaction is endothermic and equilibrium-limited, and therefore requires a continuous external heat supply. To achieve satisfactory CO selectivity, suppress methane formation via the competing Sabatier reaction, and prevent catalyst deactivation by coking, RWGS reactors are typically operated at high temperatures ($T > 850^\circ\text{C}$) [31]. Operating at elevated pressure, typically around 20 bar, can enhance reaction kinetics and reduce the need for intermediate compression steps downstream. However, higher pressures thermodynamically favour side reactions such as methanation and carbon deposition via carbon monoxide decomposition (Boudouard reaction), as these pathways result in a net decrease in the number of gas-phase molecules, requiring careful optimisation of operating conditions [31].

1.1.2.2.2 Fischer-Tropsch synthesis


The second stage converts syngas, a mixture of CO and H_2 , into hydrocarbons through a catalytic chain-growth mechanism. The resulting product distribution is inherently broad, spanning light gases, naphtha, kerosene, diesel, and heavy waxes (C_{22+}) and is commonly described by the Anderson-Schulz-Flory (ASF) distribution, which predicts a geometric decrease in the molar fraction of hydrocarbons with carbon number according to the chain-growth probability parameter α [11, 32].

The properties of the resulting syncrude are mainly determined by the catalyst type and operating conditions, which define two main process categories: low-temperature Fischer-Tropsch (LTFT, $200\text{--}257^\circ\text{C}$) and high-temperature Fischer-Tropsch (HTFT, $300\text{--}350^\circ\text{C}$), each associated with specific catalyst formulations Fe-HTFT, Fe-LTFT, or Co-LTFT [11, 33]. The carbon number distribution and typical compound classes for each configuration are summarised in Table 1.1.

Table 1.1: Carbon number distribution and typical compound classes of high-temperature Fischer-Tropsch and low-temperature Fischer-Tropsch products, excluding $\text{C}_1\text{--C}_2$ hydrocarbons [12].

Carbon number distribution (mass %)			Compound classes		
Fraction	HTFT	LTFT	Class	HTFT	LTFT
$\text{C}_3\text{--C}_4$, LPG	30	10	Paraffins	> 10%	Major product
$\text{C}_5\text{--C}_{10}$, naphtha	40	19	Olefins	Major product	> 10%
$\text{C}_{11}\text{--C}_{22}$, distillate	16	22	Aromatics	5-10%	< 1%
C_{22} and heavier	6	46	Oxygenates	5-15%	5-15%
Aqueous products	8	3	S- and N-species	None	None
			Water	Major by-product	Major by-product

For kerosene-oriented production within a PtL scheme, Co-LTFT is the most appropriate configuration. Cobalt catalysts operating at low temperature favour high α values, resulting in a syncrude heavily weighted toward the C_{22+} wax fraction, composed predominantly of n-paraffins and n-olefins, with negligible aromatics, oxygenates, and absence of sulfur or nitrogen species (see Appendix A.1 for detailed syncrude compositions) [32]. The main reaction pathways in Co-LTFT synthesis, leading respectively to alkanes and olefins, are described by Reactions (1.7) and (1.8).



The Co-LTFT syncrude can therefore be described, to a good approximation, as a binary mixture of n-paraffins and n-olefins, without the heteroatom impurities that typically complicate catalyst selection and reactor operation in conventional hydrocracking, as further discussed in Section 1.3.4. HTFT operation, by contrast, favours lighter and more olefinic products that are unsuitable for jet fuel applications [33].

Since a significant fraction of the Co-LTFT product mixture falls in the heavy C_{22+} wax range, the reactor outlet is best regarded as a synthetic crude (syncrude) rather than a finished fuel. Fischer-Tropsch synthesis should therefore be understood as the key hydrocarbon-forming step within the PtL chain, with the adjustment of product quality and the maximisation of kerosene yield delegated to downstream upgrading operations [33].

1.1.2.2.3 Upgrading

The raw FT syncrude, composed of a wide distribution of paraffinic hydrocarbons spanning from light gases to heavy waxes, cannot be used directly as aviation fuel due to specifications detailed in Section 1.2.2.1 and must therefore undergo a dedicated upgrading stage. The nature of this upgrading step, particularly the selection of hydrocracking as the central conversion technology, is examined in detail in Section 1.3.

1.1.2.2.4 Example: the ULiège process

The Chemical Engineering Department at the University of Liège (ULiège), under the leadership of Prof. Léonard, is developing an experimental pilot-scale installation for kerosene production, targeting an output of approximately 200 mL per day. The process integrates the reverse water-gas shift reaction (1.6) followed by the Fischer-Tropsch synthesis to produce C_8 – C_{16} hydrocarbons.

The broader industrial objective of the Neutral-Kero-Lime (NKL) project is to produce 250 000 m³ of kerosene per year, corresponding to approximately 40% of Liège Airport's annual fuel consumption, using 0.15 Mt of H₂ requiring roughly 1 GW of electrolysis capacity and 1 Mt of CO₂ sourced from Direct Air Capture [34].

1.1.3 Power-to-Chemicals

Power-to-Chemicals (PtC) refers to the use of renewable electricity to produce chemical compounds such as feedstocks, reagents, or intermediates, whose primary value lies in their chemical function rather than their energy content, distinguishing this branch from other PtX pathways. This approach is particularly relevant for the chemical sector, which today relies predominantly on fossil-based feedstocks such as natural gas, crude oil, and hard coal, and is therefore a significant source of greenhouse gas emissions [35]. By producing primary chemicals through electricity-derived synthesis gas and various downstream conversion technologies, PtC offers a promising pathway toward the defossilisation of material uses.

Two distinct decarbonisation logics underpin the PtC concept. For inorganic products such as ammonia, defossilisation is achieved by replacing fossil-derived hydrogen with electrolytic hydrogen from renewable electricity, as illustrated by power-to-ammonia (P2A) technology, which substitutes the conventional fossil hydrogen feed in the Haber-Bosch process [36]. For organic chemicals, however, carbon is an intrinsic

component of the product and cannot be eliminated. The objective is therefore not strict decarbonisation but defossilisation: fossil carbon is replaced by non-fossil carbon, typically supplied as captured CO_2 , while renewable hydrogen provides the required power. Through catalytic, thermochemical, or electrochemical conversion routes, these inputs can be transformed into primary chemicals such as methanol, olefins, or synthetic hydrocarbons [35], or alternatively via the electrocatalytic reduction of CO_2 into value-added molecules such as formic acid or ethylene [37].

1.1.4 Power-to-Heat

Power-to-Heat refers to the conversion of electrical energy into thermal energy for industrial or residential applications, via electric boilers or heat pumps. Beyond its direct decarbonisation potential, Power-to-Heat systems can facilitate the integration of variable renewable energy by enabling demand-side flexibility, large-scale thermal energy storage, and a reduction in renewable curtailment, while also providing grid-support services through aggregation mechanisms [38].

1.2 Aviation kerosene: production routes and fuel specifications

1.2.1 Conventional kerosene production

Kerosene is conventionally produced by petroleum refining, which relies on the separation of the many compounds present in crude oil through atmospheric and vacuum fractional distillation. In this process, the crude oil is heated and its constituent compounds vaporise at different temperatures; the resulting gases are then recondensed into liquid fractions at successive stages of the distillation column, as illustrated in Figure 1.2.

The lightest fraction, recovered at the top of the column, is naphtha, which is primarily processed into motor gasoline. The second fraction, representing approximately 33% of the crude oil input, constitutes the raw material for jet fuel production; it is further processed in a distillate hydrotreater to yield kerosene and special solvents. This is followed by the heavier gas oil fraction, which encompasses diesel fuel and heating oil. Finally, although the bottoms fractions or residual fraction can be used as heavy boiler fuel, it is usually vacuum distilled first to yield more high-value distillate [9].

Kerosene thus originates as a straight-run petroleum fraction with a boiling range of approximately $205\text{--}260^\circ\text{C}$ [39], of which only the $\text{C}_8\text{--C}_{16}$ cut is relevant for conventional jet fuel applications [9, 40], since jet fuel itself belongs to the kerosene-range fuels [11]. This contrasts sharply with Power-to-Liquids routes, where the kerosene fraction does not exist as such in any natural feedstock but must instead be synthesised from CO_2 and renewable hydrogen through a chain of catalytic conversion steps.

According to the International Energy Agency, global demand for jet fuel and kerosene reached 7 702 thousand barrels per day in 2025, representing approximately 7% of total oil demand [41], a figure that continues to rise year on year, underlining the urgency of developing lower-carbon alternatives.




Figure 1.2: The conventional route for jet fuel production from crude oil [9].

1.2.2 Sustainable aviation fuels

Sustainable Aviation Fuels (SAF) are low-carbon alternatives to conventional fossil jet fuel, produced from non-fossil feedstocks such as biomass, organic residues, or captured CO₂.

SAF production pathways can be categorised according to their feedstock origin into two main families: biofuels, derived from biomass processing, and synthetic e-fuels, produced from renewable electricity via Power-to-Liquid routes [10], as illustrated in Figure 1.3. While biofuel pathways, most notably hydroprocessed esters and fatty acids (HEFA), have already reached commercial scale, PtL-based SAF, combining electrolytic hydrogen with captured CO₂ through either the Fischer-Tropsch or methanol-to-jet pathway, remains a more recent development with a TRL of 5-6 [10, 42].




Figure 1.3: Main production routes for sustainable aviation fuel [10].

Despite SAF being recognised by the International Air Transport Association (IATA) as the most promising short-term strategy to reduce CO₂ emissions in the aviation sector, global SAF production reached only approximately 1.9 million tonnes in 2025, representing a mere 0.6% of total jet fuel [43]. Large-scale production remains hindered primarily by high production costs, estimated at 1.2 to 7 times the market price of conventional fossil jet fuel [42].

To stimulate large-scale SAF production and reduce production costs with increasing technological maturity, the European Union adopted the ReFuelEU Aviation regulation in 2023, which mandates a gradual increase in the minimum SAF share in jet fuel blends at EU airports: 2% in 2025, 6% in 2030 — of which 1.2% must be synthetic fuel — rising to 20% in 2035, with a long-term target of 70% by 2050, including 35% of synthetic fuel [21].

1.2.2.1 Quality specifications and standards

For SAF to be commercially permitted for usage, it must be certified by the ASTM D7566 Standard Specification for Aviation Turbine Fuel Containing Synthesized Hydrocarbons, set by the American Society for Testing and Materials (ASTM). The qualification process for new candidates is specified by ASTM D4054, and once a production route is qualified, it is added as an annex to ASTM D7566. These standards ensure compatibility with conventional fossil-based jet fuels (Jet A/ Jet A-1), whose properties are specified in ASTM D1655 [10].

In their current state of development, SAF must be blended with conventional fossil jet fuel, with a maximum certified blending rate of 50%. This limit reflects the very low aromatic content of currently certified synthetic paraffinic kerosene (SPK) routes (see Table 1.2), which can compromise the seal compatibility of aircraft engines [10, 44]. Research is ongoing to raise the certified blending rate towards

100% by integrating dedicated upgrading steps such as the aromatisation of light fractions discussed in Section 1.3.1 to reach the required aromatic content and therefore unlock the full decarbonisation potential of SAF.

Table 1.2 compares selected ASTM property requirements for Jet A/A-1 with those of several SAF pathways. These specifications impose key constraints on the design of an FT-based production process: the boiling-range, flash-point and freezing-point requirements together define the kerosene-range target fraction, which is commonly reported as C₉–C₁₅, C₈–C₁₆, or similar carbon-number ranges depending on jet fuel applications considered [11, 40]. In the present work, as explained Section 1.1.2.2.4, the C₈–C₁₆ definition is retained because it is consistent with the fraction of interest defined in the broader Power-to-Kerosene process developed at ULiège, into which the hydrocracking unit is intended to be integrated.

Table 1.2: Selected ASTM property requirements for Jet A/A-1 and representative SPK-type SAF pathways [10].

Property requirement	Jet A/A-1	SAF (FT-SPK/HEFA-SPK/ATJ-SPK)
Acidity (total mg KOH g ⁻¹)	< 0.1	< 0.015
Distillation, 10% recovered (°C)	< 205	< 205
Distillation, final boiling point (°C)	< 300	< 300
Flash point (°C)	> 38	> 38
Freezing point (°C)	< -40	< -40
Viscosity at -20°C (mm ² s ⁻¹)	< 8	< 8
Density at 15 °C (kg m ⁻³)	775-840	730-770
Aromatics content	8-25 vol%	< 0.5 wt%
Sulfur content (wt%)	0.3	0.0015
Trace metals*(mg kg ⁻¹ per metal)	No requirement	0.1

* Metals considered: Al, Ca, Co, Cr, Cu, Fe, K, Li, Pb, Pd, Mg, Mn, Mo, Na, Ni, P, Pt, Sn, Sr, Ti, V and Zn.

1.3 Upgrading of Fischer-Tropsch syncrude

1.3.1 Upgrading technology selection

The carbon-number-based refining framework presented throughout this section, together with the technology-selection logic of Figure 1.4, follows the analysis of de Klerk [11]. As described in Section 1.1.2.2.2, the syncrude produced by Fischer-Tropsch synthesis is composed of a broad hydrocarbon distribution theoretically spanning from C₁ to C_∞. Due to the regulatory specifications detailed in Section 1.2.2.1, these different fractions cannot be used directly and must each undergo a specific refining treatment to meet aviation fuel specifications. As shown in Figure 1.4, this refining treatment is adapted to the hydrocarbon cut being processed i.e carbon-number-based refining. It should be noted that the boundaries between fraction categories are defined slightly differently across the literature; the present work focuses on the conversion of the C₁₇+ fraction (C₁₇ and heavier), while the target kerosene fraction is defined as C₈–C₁₆.

The lightest hydrocarbons are too short to fall within the jet fuel boiling range and must therefore be lengthened or aromatized. Oligomerization can extend their chain length into the naphtha or kerosene range, while aromatisation converts alkanes and alkenes into aromatics, which are currently lacking in SAF and contribute to seal compatibility in aircraft engines. The resulting aromatic alkylation products can further be directed toward the kerosene fraction.

The naphtha fraction is similarly too light for direct use as jet fuel and can be converted into aromatics suitable for blending through catalytic reforming or co-processing with light hydrocarbons in a co-feeding or less preferred aromatisation unit.

The kerosene fraction falls within the correct boiling range but, due to the high linearity of Co-LTFT paraffins, exhibits a freezing point that is too high to meet Jet A/A-1 specifications. To address this, hydroisomerization is applied: it increases the degree of branching of linear paraffins, thereby lowering the freezing point, while simultaneously saturating any residual olefins and removing oxygenates.

Finally, the heavier fractions represent the dominant product of Co-LTFT synthesis and constitute the primary upgrading challenge. These compounds must be both cracked to shorter chain lengths and hydroisomerised to improve cold-flow properties. Hydrocracking achieves both simultaneously and is therefore the most appropriate technology for this fraction, as examined and justified in detail in the following section [11].




Figure 1.4: Technology selection for the refining of syncrude to jet fuel [11].

1.3.2 Justification for the use of hydrocracking

Hydrocracking is a catalytic process operating in the presence of hydrogen, in which acid-catalyzed cracking and (de)hydrogenation proceed simultaneously over a bifunctional catalyst [11, 33]. Therefore, hydrocracking is particularly suitable for wax upgrading, since it adjusts the carbon number distribution through cracking while introducing branching through hydroisomerization, which helps improve the cold-flow properties of the resulting jet fuel fraction [33, 45]. To date, it remains the only heavy-end upgrading technology capable of delivering a good quality distillate as its main product [11]. The underlying reaction mechanism, along with catalyst selection, operating conditions, kinetics and reactor configurations, are discussed in Section 1.4.

1.3.3 Possible cracking routes

In the literature, hydrocracking of LTFT waxes is consistently presented as the most suitable upgrading technology for heavy fractions to maximise the overall middle distillate yield [6, 11, 32, 46–48]. However, it is important to understand why and against which alternatives this conclusion holds. Cracking technologies can be divided into three main categories: thermal cracking, fluid catalytic cracking (FCC), and hydrocracking. The advantages and limitations of each route must therefore be evaluated in the specific context of Co-LTFT wax upgrading to jet fuel.

1.3.3.1 Thermal cracking

Thermal cracking is the earliest refining technology, defined as the thermal decomposition of heavy hydrocarbon fractions under high temperature (324–500 °C) and pressure, without any catalyst, to produce lighter molecules [11, 33, 49]. The process relies solely on heat to provide sufficient energy for C–C bond rupture, without consuming hydrogen as a reactant, though hydrogen or steam may be introduced as a stripping gas to limit the formation of light hydrocarbon gases.

Its main advantages are operational simplicity, since no catalyst is required, and a naturally low light-gas yield, since the cracking rate decreases with decreasing chain length, thereby limiting overcracking

of lighter fractions. However, thermal cracking suffers from several limitations. It yields a broad, bimodal carbon number distribution and produces a distillate fraction rich in olefins ($\sim 40\%$ mass) [49]. Furthermore, the absence of a hydrogenation function means these unsaturated species cannot be saturated in situ. Additionally, inorganic species carried over from the upstream FT catalyst (e.g. cobalt fines and silica from the support) progressively accumulate on reactor internals during operation, leading to fouling and flow obstruction [49]. For jet fuel production from Co-LTFT waxes, thermal cracking is therefore unsuitable: its high olefin yield is incompatible with SAF specifications, and direct comparison with hydrocracking at similar conversion levels shows that it achieves a lower distillate yield while requiring operating temperatures roughly $70\text{ }^{\circ}\text{C}$ higher [49]. For these reasons, thermal cracking is better suited to the production of chemicals such as ethylene, candle wax and n - α -olefins [6, 11].

1.3.3.2 Fluid catalytic cracking (FCC)

Fluid catalytic cracking (FCC) is one of the most important processes in petroleum refining history, defined as an acid-catalyzed cracking process operating in the absence of hydrogen, at temperatures of $480\text{--}550\text{ }^{\circ}\text{C}$, pressures of $0.1\text{--}0.3\text{ MPa}$, and contact times below 10 seconds [11]. The feed is brought into contact with a hot zeolitic acid catalyst, which initiates cracking via carbenium ion intermediates. The coke formed on the catalyst during cracking is combusted in a regenerator to supply the heat driving the process.

Originally developed for the conversion of heavy petroleum-derived residues, FCC is designed to crack carbon-rich, hydrogen-poor fractions into lighter and more valuable products, primarily gasoline and light cycle oil, by rejecting the excess carbon as coke on the catalyst [11]. However, FCC presents a poor technology fit with Co-LTFT syncrude [50]. The process was conceived as a carbon rejection technology, designed for feedstocks with a low H:C ratio. Co-LTFT syncrude, by contrast, is composed predominantly of linear paraffins with a higher H:C ratio; there is simply no carbon excess to reject. As a direct consequence, insufficient coke forms on the catalyst, compromising the heat balance of the regenerator and requiring additional fuel to sustain operation [50, 51]. Furthermore, the absence of hydrogen results in a product slate rich in unsaturated molecules, which is incompatible with jet fuel specifications [52]. Producing jet fuel from FCC products would require additional downstream steps, oligomerization and hydrogenation, making the overall pathway less efficient than hydrocracking [11]. FCC may find application in specific contexts where gasoline production is prioritised, but the limitations outlined above make it unsuitable as a primary upgrading route for Co-LTFT waxes into jet fuel, and motivate the selection of hydrocracking instead [11, 33].

1.3.4 Comparison with conventional hydrocracking

The fundamental differences between Fischer–Tropsch syncrude hydrocracking and conventional hydrocracking of petroleum-derived fractions arise primarily from the feedstock composition. As described in Section 1.1.2.2.2, Co-LTFT syncrude is dominated by linear n -paraffins with minor amounts of olefins, contains only very low amounts of aromatics and cyclic hydrocarbons, and is virtually free of sulfur- and nitrogen-containing compounds. Petroleum-derived heavy fractions, by contrast, are compositionally far more complex, containing alkanes, cycloalkanes, aromatics, and a variety of heteroatomic species bearing sulfur, nitrogen, and oxygen [53]. These compositional differences have direct and significant consequences on operating conditions, catalyst selection, and overall process design.

In conventional hydrocracking, a substantial fraction of the hydrogen consumed is spent on aromatic saturation and heteroatom removal, reactions that require high pressures, typically above 10 MPa [12]. In the case of LTFT wax, the absence of aromatics and heteroatoms removes this requirement, so that hydrogen is mainly consumed by olefin saturation and hydrocracking reactions. As a result, catalyst coking is suppressed and the process can be operated at significantly lower pressures, as summarised in Table 1.3 [12]. A further consequence of the feedstock composition is the thermal behaviour of the reactor: the endothermic cracking reactions and the exothermic olefin saturation reactions largely compensate

one another, rendering the process nearly isothermal, in contrast to the clearly exothermic character of conventional hydrocracking [12].

Table 1.3: Typical conditions for conventional and LTFT wax hydrocracking [11, 13–16].

	Conventional hydrocracking	LTFT hydrocracking
Operating conditions		
Temperature ($^{\circ}\text{C}$)	350-430	325-375
Pressure (MPa)	10-20	3.5-7
Space velocity (h^{-1})	0.2-2	0.5-3
H_2 :feed ratio ($\text{m}^3 \text{m}^{-3}$)	800-2000	500-1800
Conversion properties		
Cracking conversion (%)	70-100	20-100
H_2 consumption (mass% of feed)	1.4-4	< 1
Heat release	Exothermic	Near-isothermal
Catalyst characteristics		
Hydrogenation function	Sulfided NiMo or NiW	Pt, Pd or non-sulfided Ni-based systems
Acidic support	Amorphous silica-alumina or zeolite	Silica-alumina or zeolite

The sulfur-free nature of Co-LTFT syncrude also has direct implications for catalyst selection. In conventional hydrocracking, the sulfur present in the feed poisons noble-metal catalysts and imposes the use of sulfided base-metal systems, typically Ni/Mo or Ni/W on silica-alumina or zeolite supports, which are resistant to poisoning but comparatively less active and selective. With Co-LTFT wax this constraint disappears, and unsulfided Pt- or Pd-based catalysts can be employed directly: combined with the high intrinsic reactivity of heavy n-paraffins, their higher activity and selectivity enable high conversion under operating conditions significantly milder than those required in conventional petroleum hydrocracking [12, 32].

It is worth noting that conventional hydrocracking technology could, in principle, be applied to LTFT wax [11]. However, doing so would result in several inefficiencies. Operating conditions would be unnecessarily severe, as they are designed to handle aromatic saturation and heteroatom removal, constraints that are simply absent in LTFT wax [11]. More critically, sulfided base-metal catalysts require a minimum H_2S level to maintain their active sulfided phase during hydrocracking of sulfur-free LTFT wax. In the absence of sulfur addition, catalyst activity cannot be sustained, while changes in the metal/metal-sulfide equilibrium alter the metal-acid balance, leading to deactivation and changes in product selectivity [54]. At the same time, the inherent advantages of noble-metal catalysts, their high activity, selectivity, and compatibility with a linear paraffinic feedstock, would not be exploited. For these reasons, the hydrocracking of LTFT syncrude is best regarded as a dedicated process, distinct from conventional petroleum hydrocracking, and should be designed accordingly.

1.4 Hydrocracking of LTFT syncrude

1.4.1 Reaction mechanism and catalysts

Hydrocracking of Fischer-Tropsch waxes is carried out over a bifunctional catalyst, which combines two distinct types of active sites involved in a sequential reaction pathway. The reaction proceeds in two coupled steps: a reversible hydroisomerization that converts linear n-paraffins into their branched iso-paraffin counterparts, followed by the cracking of these iso-paraffins into shorter paraffinic fragments. These two steps are carried by metallic sites responsible for hydrogenation and dehydrogenation (HD/DHD), and Brønsted acidic sites responsible for isomerization and carbon-carbon bond cleavage [4, 45, 55]. This dual functionality is central to the bifunctional hydrocracking mechanism, in which long-chain n-paraffins, constituting the dominant fraction of Co-LTFT syncrude, are highly saturated molecules and are therefore too stable to be directly activated on an acidic site alone [32]. The metallic function is required to first generate reactive olefinic intermediates, which are subsequently processed by the acid function. It should be noted that any olefins already present in the feedstock, a non-negligible fraction in Co-LTFT syncrude,


are considered to be instantaneously hydrogenated to the corresponding n-paraffins upon entering the reactor, given their considerably higher hydrogenation reactivity than n-paraffins [5, 55]. These n-paraffins then enter the bifunctional conversion sequence described below.

1.4.1.1 Bifunctional reaction pathway

The most widely accepted reaction mechanism for hydrocracking of n-paraffins over bifunctional catalysts, established since the seminal contributions of Weisz [56] and Coonradt and Garwood [57], and thoroughly investigated by Weitkamp, Froment and co-workers, involves eight elementary steps that can be grouped into three stages [4, 6, 32]:

1. **Activation on the metallic function:** Upon adsorption on a HD/DHD metallic site, the n-paraffin is dehydrogenated to the corresponding olefin, which then diffuses toward an acidic site. For this step to proceed, the presence of hydrogen gas is required [32, 45].
2. **Processing on the acidic function:** The n-olefin is protonated on a Brønsted acidic site to form a carbenium ion. This carbenium ion can then undergo isomerization on the acid function through hydride shifts, methyl shifts, and protonated cyclopropane (PCP)-mediated branching. While hydride and methyl shifts redistribute the charge and carbon skeleton, PCP branching progressively increases the degree of substitution, forming mono- and then multibranched iso-carbenium ions [6]. Cracking via β -scission only becomes accessible once sufficient branching has been introduced, since only di- and multibranched isoalkanes can crack according to energetically favourable β -scission modes; direct cracking of n-paraffins or monobranched species is slow and can be neglected in practice [55, 58]. β -scission proceeds by rupture of the C-C bond in the β position relative to the positively charged carbon, yielding a lighter carbenium ion and a lighter olefin [45].
3. **Stabilisation on the metallic function:** After deprotonation on the acid site, the olefinic products diffuse back to a HD/DHD metallic site, where they are hydrogenated to the corresponding paraffins, either isoparaffins (isomerization products) or shorter n-paraffins (cracking products). The stabilized products are then desorbed from the catalyst [4, 45].

The overall sequence can be summarised as follows [59]:



It should further be noted that hydrocracking of Fischer-Tropsch waxes is a catalytic surface reaction taking place in a three-phase environment, in which reactants are distributed between a vapour phase and a liquid phase in contact with the solid catalyst [4]. The distribution of components between the two phases, governed by vapour-liquid equilibrium, is affected by operating conditions and directly influences the local concentrations of reactants at the catalyst surface. This aspect has important consequences for kinetic modelling and reactor design, and will be examined further in Sections 1.4.3 and 1.4.5.

1.4.1.2 Ideal and non-ideal bifunctional catalysts

The efficiency of the bifunctional mechanism depends primarily on the balance between the two catalytic functions. In the terminology introduced by Weitkamp [60], an ideal bifunctional hydrocracking catalyst is one in which the hydrogenation/dehydrogenation reactions proceed at quasi-equilibrium, such that the rate-limiting step is located exclusively on the acidic function [32]. This condition is achieved when:

- The two functions are in close spatial proximity, satisfying the *Weisz intimacy criterion*, so that the transport of olefinic intermediates between metallic and acidic sites does not limit the overall reaction rate [32].

- The HD/DHD function is sufficiently strong to maintain the maximum thermodynamically achievable concentration of intermediate olefins at the acid sites, and to hydrogenate cracked olefinic fragments rapidly before they undergo further acid-catalyzed reactions [32].

In such an ideal catalyst, three key features follow [32]:

1. The process can operate at relatively low temperatures.
2. High selectivities for hydroisomerization are achievable.
3. Pure primary cracking occurs: each heavy molecule is cracked once, and the resulting fragments leave the catalyst before undergoing further cracking. The resulting carbon number distribution of cracked products is therefore fully symmetrical and centred at approximately half the carbon number of the parent molecule [32].

By contrast, a non-ideal catalyst, one in which the HD/DHD function is insufficiently strong relative to the acidic function, allows cracked intermediates to remain in the vicinity of acidic sites long enough to undergo secondary cracking. This shifts the product distribution toward lighter species (typically C_3 - C_5) and significantly reduces the selectivity for middle distillates [32]. Following Marcilly's definition, the ideal hydrocracking catalyst is the one for which primary cracking is maintained up to the highest possible conversion [61]. It is worth noting that industrial catalysts, provided they present a well-balanced ratio between acidic and metallic sites, generally exhibit behaviour comparable to the ideal case [62].

A good hydrocracking catalyst for LTFT wax upgrading is therefore not necessarily the most acidic or the most active one, but the one that achieves the best compromise between wax conversion and middle distillate selectivity [32]. In the context of the present work, this translates into a clear design requirement: maximising the C_8 - C_{16} yield demands a catalyst in which the metallic function is strong enough to rapidly hydrogenate primary cracking products, thereby preventing both secondary cracking and overcracking of the middle distillates already present in the feedstock [17].

1.4.1.3 Catalyst materials

The selection of both the metallic component and the acidic support governs the overall activity, selectivity, and resistance to deactivation of the hydrocracking catalyst.

1.4.1.3.1 Acidic supports

The acidic function is provided by the support material, for which three main families are encountered in the hydrocracking literature [4, 55]: amorphous oxides or mixtures of oxides (e.g., HF-treated Al_2O_3 , SiO_2 - Al_2O_3); crystalline zeolites (Y, Beta, Mordenite); and silicoalumina-phosphates (SAPO-11, SAPO-31, SAPO-41) [5]. Notably, studies on both zeolitic and amorphous catalysts have consistently concluded that the underlying reaction mechanism is the same regardless of the nature of the acidic support [6], with carbenium ion chemistry governing the isomerization and cracking steps in both cases.

1.4.1.3.2 Metallic components

The most commonly used metallic systems are platinum (Pt), palladium (Pd), or bimetallic sulfided systems such as Ni/Mo, Ni/W, and Co/Mo [4, 45, 55].

Catalysts based on noble metals, particularly Pt, exhibit superior performance in terms of selectivity for hydroisomerization and product distribution compared to non-noble metal catalysts [4, 45, 63, 64]. This advantage arises from their high intrinsic hydrogenation/dehydrogenation activity, which facilitates the ideal catalyst behaviour described above. However, as discussed in Section 1.3.4, noble metals are susceptible to sulfur poisoning, which makes them incompatible with conventional petroleum-derived feedstocks. In the specific case of Co-LTFT syncrude, which is virtually sulfur-free, this limitation does not apply, and

Pt- or Pd-based catalysts can be used without risk of deactivation by sulfur. This represents one of the key advantages of hydrocracking LTFT waxes over conventional feedstocks.

Non-noble metal systems, generally used in sulfided form, are the basis for conventional hydrocracking, but they become less suitable for sulfur-free FT feeds, as the absence of sulfur prevents the maintenance of the active sulfided phase and disrupts the metal-acid balance, leading to changes in selectivity and gradual deactivation [54]. Furthermore, certain reduced non-noble metals such as Ni and Co can promote deep hydrogenolysis, resulting in excessive formation of light gases (C_1 - C_2), which is detrimental to the distillate yield [32].

Although Co-LTFT syncrude is essentially free of sulfur and nitrogen, it may contain oxygenated compounds such as alcohols, acids, ketones, and esters. These species can affect both catalytic functions: alcohol-derived water may inhibit acid sites and reduce activity, whereas carboxylic acids or their decomposition products may alter the HD/DHD function and disturb the metal-acid balance. Catalyst selection for LTFT wax hydrocracking must therefore account not only for sulfur tolerance, but also for the possible influence of oxygenates on activity, selectivity, and stability [32], although normally their presence is very limited.

1.4.1.3.3 Resistance to deactivation


A further advantage of hydrocracking over purely acid-catalyzed cracking technologies is its superior resistance to catalyst deactivation. Because coke precursors are continuously hydrogenated by the metallic function, coke formation is strongly reduced, and the catalyst lifetime can extend to years of continuous operation, in contrast to fluid catalytic cracking catalysts, which must be regenerated on timescales of seconds [11].

1.4.2 Stoichiometry

The reaction mechanism described in Section 1.4.1 must be expressed as global stoichiometric relations before it can be implemented in a reactor model. This formulation does not address the rate at which reactions occur, but rather defines how the molecular structure and carbon-number distribution of the reacting paraffins are modified within the overall hydrocracking network. Two stoichiometric transformations must therefore be distinguished. Hydroisomerization preserves the carbon number and relates a linear paraffin to its branched isomers, whereas the cracking step changes the carbon-number distribution by splitting a branched paraffin into two shorter paraffinic fragments. These two transformations are part of the same bifunctional hydrocracking process, but they play different roles in the reaction network and are therefore introduced separately below.

1.4.2.1 Hydroisomerization stoichiometry

Hydroisomerization is the overall transformation in which a linear paraffin is converted to one of its branched isomers while preserving the carbon number. As discussed in Section 1.4.1, the reaction proceeds through successive dehydrogenation, protonation, skeletal rearrangement of the carbenium ion on the acid site, deprotonation, and hydrogenation. Since these two hydrogen-transfer steps compensate each other in the overall balance, hydroisomerization does not involve any net hydrogen consumption. Its global stoichiometry can therefore be written as Reaction (1.10).



In this expression, $\text{iso-}C_nH_{2n+2}$ does not denote a single molecular species, but a pseudo-component grouping all branched isomers with the same carbon number n . This representation is commonly adopted in Fischer–Tropsch wax hydrocracking models because the number of possible skeletal isomers increases rapidly with chain length, and because chromatographic analysis does not generally allow mono-, di-, and multibranched isomers to be distinguished for heavy hydrocarbons. For instance, Pellegrini et al., report


that decane ($C_{10}H_{22}$) already has 75 possible isomers, while eicosane ($C_{20}H_{42}$) has more than 3.6×10^5 possible isomers [5]. For this reason, the models developed in the series of papers by Pellegrini et al. retain the simplified distinction between n-paraffins and lumped iso-paraffins, rather than tracking each individual branched molecule separately [3–7].

The reversible character of Reaction (1.10) is important because hydroisomerization is not represented as a one-way conversion towards branched products, but as an equilibrium between the linear paraffin and the set of branched isomers of the same carbon number. Since $iso-C_nH_{2n+2}$ represents an ensemble of branched structures rather than a single molecular species, the corresponding equilibrium constant is a lumped equilibrium constant. Pellegrini et al. emphasise that this constant cannot be inferred from the thermodynamics of one representative isomer only. Indeed, $iso-C_nH_{2n+2}$ is a lumped pseudo-component whose equilibrium concentration results from the sum of all branched isomers with carbon number n . Since the number of possible branched isomers increases rapidly with chain length, this multiplicity must be included when evaluating the lumped equilibrium constant [58].

1.4.2.2 Cracking step stoichiometry

The stoichiometric description of hydrocracking stems directly from the β -scission step introduced in Section 1.4.1. In bifunctional hydrocracking, C-C bond cleavage does not occur on the saturated n-paraffin itself, but on a carbenium ion formed after dehydrogenation, protonation, and skeletal rearrangement. The β -scission of this carbenium ion produces two carbon fragments, which are subsequently converted, after deprotonation and hydrogenation on the metallic function, into saturated paraffinic products. This sequence explains why primary hydrocracking leads to two fragments as products [32].

For an iso-paraffin of carbon number n , a primary cracking event leading to two fragments of carbon numbers j and $n - j$ is represented by Reaction (1.11).



In Reaction (1.11), the notation C_jH_{2j+2} only specifies the carbon number and the saturated paraffinic nature of the fragment. It does not imply that the product is necessarily a linear paraffin. Depending on the structure of the reacting carbenium ion and on the cracking mode involved, the hydrogenated fragments may be either a n-paraffin or an iso-paraffin [6]. The stoichiometric equation therefore defines the carbon-number split, whereas the distinction between normal and branched fragments will be specified in the kinetic model.

The single mole of hydrogen appearing in Reaction (1.11) corresponds to the hydrogen required to saturate the two fragments generated by the C-C bond rupture. The global stoichiometry, however, does not specify which C-C bond is cleaved, and therefore does not determine the value of j . This additional information is provided by the product-distribution assumption adopted in the model. In the following, two such descriptions are considered: the midpoint approximation, in which only the central cleavage is retained, and the breakage probability function, in which several cleavage positions are allowed and weighted statistically.

A central experimental observation underlying both approaches is that primary hydrocracking over ideal bifunctional catalysts tends to produce a symmetric carbon-number distribution of fragments [32]. In this context, symmetry should not be understood as a molecule-by-molecule cleavage exactly at the centre of the chain. Rather, it means that the formation of a fragment C_j is associated with the formation of its complementary fragment C_{n-j} , so that the overall product distribution is mirrored around half of the parent carbon number. For example, primary cracking of a C_{12} paraffin may contribute to the pairs C_3/C_9 , C_4/C_8 , C_5/C_7 , and C_6/C_6 , leading to a distribution centred around C_6 without requiring every molecule to break exactly into two equal halves. Such behaviour has been reported for several model n-alkanes, also referred to as normal paraffins, over bifunctional catalysts, provided that secondary cracking remains limited [6, 65]. When the primary fragments undergo further cracking, the distribution progressively loses this symmetry and shifts towards lighter hydrocarbons [6, 32].


1.4.2.2.1 Midpoint approximation

The simplest way to represent the symmetric character of primary hydrocracking in a reactor model is to retain only the central cleavage of the carbon chain. In this approximation, the full distribution of possible fragment pairs is collapsed into a single pair of products located at, or as close as possible to, half of the parent carbon number. The midpoint approximation therefore preserves the two-fragment nature of primary hydrocracking, but removes the explicit summation over all possible C-C bond cleavage positions.

For an even carbon number n , as illustrated in Figure 1.5a, the central cleavage gives two fragments of identical size, as shown in Reaction (1.12).



For an odd carbon number, exact division into two equal fragments is not possible. The molecule is therefore split into the two neighbouring carbon numbers, as represented in Reaction (1.13).



This simplifying assumption was adopted in the early Politecnico di Milano models [3–5], based on earlier experimental observations by Froment and co-workers showing a symmetric cracking distribution with a pronounced maximum near the middle of the carbon chain [5,65]. The resulting stoichiometry is particularly convenient because each reacting iso-paraffin is associated with only one product pair, which greatly simplifies the component material balances.

The midpoint approximation should therefore not be interpreted as a detailed description of the real molecular cracking spectrum. It is a single-pair representation of a broader symmetric distribution. Its main limitation is that it concentrates all primary products around $n/2$, whereas experimental product distributions extend over several carbon numbers. This limitation motivated the subsequent introduction of a breakage probability function, in which several cleavage positions are retained and weighted statistically.

1.4.2.2.2 Breakage probability function


The midpoint approximation provides a simple stoichiometric representation of symmetric primary cracking, but it collapses the whole product spectrum into a single pair of fragments. This becomes restrictive when the objective is to reproduce the detailed carbon-number distribution of the products. For this reason, Gamba et al. introduced a breakage probability function in which the central cleavage assumption is replaced by a statistical distribution of possible C-C bond ruptures [6,7]. This refinement was designed to correct the product distributions predicted by the previous model, while preserving the two-fragment nature and the overall symmetry of primary hydrocracking.

The probability function is not intended to provide a detailed elementary description of every possible β -scission pathway. Its role is to translate the experimental trends observed for primary hydrocracking of model paraffins into a stoichiometric distribution rule for the cracking products [7]. These distributions are generally centred around half of the parent carbon number, show negligible formation of C_1 and C_2 fragments, and display lower amounts of C_3 and C_{n-3} compared with the intermediate fragments. For this reason, Gamba et al. assigned a zero probability to cleavages leading to C_1 , C_2 , C_{n-1} , and C_{n-2} , a half-probability to cleavages leading to C_3 and C_{n-3} , and an equal probability to the intermediate fragments from C_4 to C_{n-4} [6,7]. If A is the probability assigned to each of the $n-7$ intermediate cleavage positions, the two quasi-terminal positions leading to C_3 and C_{n-3} are assigned a probability $A/2$. Normalizing the probabilities over all admissible cleavage positions gives Eq. (1.14).

$$(n-7)A + 2 \frac{A}{2} = 1 \quad \implies \quad A = \frac{1}{n-6} \quad (1.14)$$

On a molar-average basis, the corresponding stoichiometric distribution for the primary cracking of $iso-C_nH_{2n+2}$ is given by Reaction (1.15) and an example is shown in Figure 1.5b. Starting from

$iso-C_nH_{2n+2}$, it will break giving rise to $2An_0$ moles of each product between C_4 and C_{n-4} and to $2(A/2)n_0 = An_0$ moles of C_3 and C_{n-3} , where n_0 is the initial value of the paraffin moles [6].



Reaction (1.15) does not describe the products formed from one isolated molecule. It is an average stoichiometric representation of many primary cracking events. Each individual event still produces only two fragments, but the fractional coefficients describe how these fragments are statistically distributed over the different carbon numbers. In this sense, the breakage probability function preserves the symmetric character of primary hydrocracking, but distributes the products over several carbon numbers instead of concentrating them entirely at $n/2$.

For light hydrocarbons with $n < 8$, the general probability function is no longer directly applicable and specific rules must be introduced. In the model of Gamba et al., for example, the five C-C bonds of C_6 are assumed to break with equal probability, so that the cracking products are distributed uniformly over C_1 , C_2 , C_3 , $n-C_4$, and $n-C_5$ [6]. This rule is introduced because the general probability function assigns zero probability to C_1 and C_2 fragments, whereas methane and ethane are nevertheless detected experimentally. It is finally important to emphasise that A is not a kinetic parameter. It specifies how the two fragments are distributed in carbon number once a cracking event has occurred; the rate at which such events occur belongs to the kinetic model discussed in the following section.




Figure 1.5: Assumed product distribution for the hydrocracking of $C_{60}H_{122}$.

1.4.3 Kinetics

1.4.3.1 Overview of kinetic modelling approaches

Given the mechanistic complexity of hydrocracking described in Section 1.4.1 and the broad composition of the Fischer-Tropsch wax feedstock, developing a kinetic model for this system is particularly challenging. Simple rate expressions could in principle be used, but the objective of the present work is to develop a model that closely reflects the physical and chemical reality of the system, capturing the effect of operating conditions, the adsorption of species on the catalyst surface, and the dependence of reaction rates on chain length. To address this, two main classes of kinetic models have been proposed in the literature, single-event kinetics and pathway-level Langmuir-Hinshelwood-Hougen-Watson (LHHW) models, differing in the level of molecular detail they retain and in the assumptions they make to limit the number of adjustable parameters.

Single-event kinetic models

Single-event kinetics, developed from the 1980s onwards by G. F. Froment and co-workers [65–68], is based on a network of elementary reactions taking place on the surface of the catalyst, designed to account for the formation of each individual component of the product mixture. The key advantage is that the corresponding rate coefficients are, by construction, independent of the feedstock composition: parameters estimated from a given experimental dataset can be used to predict conversion and product distribution for feeds of very different composition [5, 65]. However, the number of elementary steps grows rapidly with chain length. For example, for n-hexadecane alone, the reaction network already requires 465 species and 1503 elementary reactions [6, 69]. Applying this framework to a full Fischer-Tropsch wax mixture up to C_{70} would therefore generate a prohibitively large network, requiring a parameter estimation effort difficult to reconcile with the experimental datasets available. Although relumping strategies have been proposed to reduce this complexity, published applications to Fischer-Tropsch wax mixtures remain limited, with examples not extending beyond C_{18} feeds [5, 70].

Pathway-level LHHW models

The LHHW formalism provides rate expressions rooted in the bifunctional mechanism of Section 1.4.1, while remaining numerically manageable for a feedstock composed of a large number of components. First applied to paraffin hydrocracking by Steijns and Froment [71] and extended by Froment in 1987 [65], it relies on the kinetic scheme (1.16).



Direct cracking of n-paraffins is neglected, as it is negligibly slow compared with the cracking of the branched isomers; cracking therefore proceeds from the isomerized species only [4]. Physical adsorption is described by the Langmuir isotherm, shown by Steijns and Froment to outperform rival isotherms in regression against experimental data [71]. The rate-determining step is the isomerization of the carbenium ion on the acid function, all other elementary steps being assumed at quasi-equilibrium [65]. The competitive adsorption of all species on the catalyst surface is represented by the common adsorption denominator defined in Equation (1.17)

$$ADSORB = \phi_{H_2} \left[1 + \sum_{i=1}^{N_C} K_{L,n}(i) \phi_n(i) + \sum_{i=4}^{N_C} K_{L,iso}(i) \phi_{iso}(i) \right] \quad (1.17)$$

where ϕ denotes partial pressure in models assuming a single vapour phase, and fugacity in models accounting for vapour-liquid equilibrium, as discussed in Section 1.4.3.3. The Langmuir constants $K_{L,n}(i)$ and $K_{L,iso}(i)$ increase with chain length, so that heavier paraffins occupy the active sites preferentially and react faster in a competitive mixture. The explicit presence of ϕ_{H_2} captures the experimentally observed inhibition of the overall rate at high hydrogen pressure (see Section 1.4.4) [65]. The rate expressions for hydroisomerization and cracking are expressed by Equations (1.18) and (1.19) respectively.

$$r_{isom}(i) = \frac{k_{isom}(i) \left(\phi_n(i) - \frac{\phi_{iso}(i)}{K_{eq}(i)} \right)}{ADSORB} \quad (1.18)$$

$$r_{cr}(i) = \frac{k_{cr}(i) \phi_{iso}(i)}{ADSORB} \quad (1.19)$$

The isomerization rate is reversible, with $K_{eq}(i)$ as the lumped equilibrium constant between the n-paraffin and all its branched isomers at carbon number i ; its value increases with chain length as the number of possible isomers grows [5]. Cracking acts irreversibly on the iso-paraffin, with both $k_{isom}(i)$ and $k_{cr}(i)$ following an Arrhenius temperature dependence. Because each parameter has a direct physical counterpart in the reaction mechanism, the model behaviour under varying operating conditions remains physically consistent.

1.4.3.2 The Pellegrini-Gambaro research series: common framework

The five papers reviewed in this section [3–7], published between 2004 and 2011 by Pellegrini, Gamba, Gambaro and their collaborators at the Politecnico di Milano and Eni, form a deliberate incremental series in which each article removes one simplifying hypothesis from its predecessor, constituting the most complete body of work on kinetic modelling of Co-LTFT wax hydrocracking to date. Beyond this progressive development, the series has also become a widely used reference in simulation studies of Fischer-Tropsch wax hydrocracking, with several subsequent works adopting the kinetic framework or close variants thereof [72–74].

Experimental basis

All five models were calibrated against the same experimental dataset, collected on a bench-scale down-flow trickle-bed reactor loaded with 9 g of powdered catalyst (average particle size: 0.625 mm) to approximate plug-flow behaviour. The catalyst is a bifunctional system composed of platinum (0.6 wt%) supported on amorphous silica-alumina. The feedstock is a C_4 - C_{70} mixture of n-paraffins. Operating conditions were explored according to a central composite design (a second-order factorial design), spanning: temperature 343–375 °C; pressure 35–60 bar; H_2 /wax mass ratio 0.06–0.15 kg/kg; and WHSV 1–3 $kg_{n-C_{in}} kg_{cat}^{-1} h^{-1}$.

Reaction scheme and kinetic equations

Every model in the series is built on the simplified kinetic scheme of Reaction (1.20), derived from the full Steijns-Froment network [65] by merging all branched isomers of carbon number i into a single lumped pseudo-component iso- C_i :



All models employ the LHHW rate expressions of Equations 1.18–1.19, involving the parameters defined in Table 2.2: $k_{isom}(i)$, $k_{cr}(i)$, $K_{eq}(i)$, $K_{L,n}(i)$, and $K_{L,iso}(i)$, with Arrhenius temperature dependence for $k_{isom}(i)$ and $k_{cr}(i)$. Parameters are estimated by non-linear regression against the experimental dataset in all articles.

Reactor model

The reactor is modelled as an isothermal plug-flow reactor, described by Equation 1.21, without axial or radial diffusion, consistently with the bench-scale trickle-bed reactor operated in downflow mode in the experimental study [6].

$$\frac{dY_i}{d\tau} = r_i(Y, T, P, \dots) \quad (1.21)$$

Y_i denotes the molar amount of species i per kilogram of inlet n-paraffins, τ is the contact time ($kg_{cat} \cdot h / kg_{n-C_{in}}$), and r_i is the net reaction rate of species i , expressed in $kmol_i / kg_{cat} / h$. Depending on the species considered, r_i includes contributions from hydroisomerization, cracking, and production by the cracking of heavier hydrocarbons.

The use of a powdered catalyst in the experimental reactor justifies neglecting inter- and intraphase transport resistances [6].

1.4.3.3 Article-by-article development

1.4.3.3.1 Pellegrini et al. (2004): The lumped model [5]

Pellegrini et al. (2004) is the founding contribution of the series, developing a LHHW-based kinetic model for Co-LTFT wax hydrocracking, with the objective of describing the behaviour of the petroleum cuts of interest (naphtha, kerosene, gas oil, and atmospheric residue) under varying operating conditions. Earlier lumped models were judged inadequate for lacking Arrhenius-type temperature dependence or the ability

to describe operating conditions [75]. The feedstock is divided into nine lumps (five n-paraffin and four iso-paraffin groups), each assigned a single constant value for all kinetic parameters, assuming a uniform species distribution within the lump. Cracking follows the midpoint approximation, all cracking products are taken to be iso-paraffins, and the reactor is treated as a single vapour phase.

The model reproduces the effects of temperature, pressure, and WHSV satisfactorily against 25 experimental tests, with correlation indices close to unity for gas oil and kerosene fractions. However, it shows no sensitivity to the H_2 /wax ratio, a deficiency the authors attribute to the absence of vapour-liquid equilibrium calculations. This paper establishes the LHHW kinetic scheme and the isothermal plug-flow reactor model that all subsequent articles in the series inherit, while the lump-based material balances, vapour-only assumption, and midpoint cracking hypothesis define the three limitations addressed in the following works.

1.4.3.3.2 Pellegrini et al. (2007): The all-components model [3]

Pellegrini et al. (2007) replaces the lump-level material balances of the 2004 model with individual component balances for all 138 species (70 n-paraffins, 67 iso-paraffins, and H_2), removing the implicit assumption of a uniform species distribution within each lump. The same LHHW rate expressions and five lump values for kinetic constants are retained. In a second stage, each kinetic constant is expressed as a continuous function of carbon number using power-law, logarithmic, or exponential forms, thereby eliminating the discontinuities at lump boundaries. The midpoint approximation and vapour-only assumption are retained.

The all-components approach produces smooth carbon-number profiles in good agreement with experimental data for n-paraffins. However, iso-paraffin concentrations in the C_5 - C_{15} range are systematically underestimated: by concentrating all cracking products near $n_C/2$, the midpoint approximation fails to populate the lighter iso-paraffin fractions. This paper introduces the component-level material balance structure and continuous kinetic constant correlations used in all subsequent models, and identifies the midpoint approximation as the next hypothesis to address.

1.4.3.3.3 Pellegrini et al. (2008): Introduction of vapour-liquid equilibrium [4]

Pellegrini et al. (2008) incorporates vapour-liquid equilibrium (VLE) calculations into the reactor model to address the insensitivity of the previous models to the H_2 /wax ratio, a variable that acts primarily through the phase distribution of hydrocarbons rather than through kinetics directly. Since critical property data for heavy paraffins are unavailable beyond approximately C_{30} in standard databases, the authors extend the method of Soave [76] to estimate T_c and P_c for n- and iso-paraffins up to C_{100} . Flash calculations are performed at each integration step using the Soave-Redlich-Kwong (SRK) equation of state with Tsonopoulos-Heidman binary interaction parameters. The resulting fugacities replace partial pressures in the LHHW rate equations, maintaining a single set of material balance equations across both phases. The midpoint approximation is retained.

The model correctly reproduces the effect of the H_2 /wax ratio on conversion: hydrogen acts as a stripping agent, shifting heavy hydrocarbons toward the vapour phase and increasing their fugacity in the rate expressions and thereby promoting their conversion. Parity plots for all four operating variables show improved alignment with the experimental diagonal compared to the previous model. The remaining discrepancies in the iso-paraffin distribution confirm the midpoint approximation as the next hypothesis to address. This paper establishes fugacity-based VLE kinetics as the thermodynamic framework for all subsequent models in the series, and provides the critical-property estimation procedures required for paraffins beyond C_{30} , enabling the treatment of the full C_4 - C_{70} feedstock.

1.4.3.3.4 Gamba et al. (2009): Introduction of a breakage probability function [6]

Gamba et al. (2009) replaces the previous midpoint cracking assumption with the breakage probability function detailed in Section 1.4.2.2.2, derived from experimental hydrocracking data available in the lit-

erature. Instead of forcing iso-paraffins to crack at the middle of the carbon chain, the new formulation distributes the cracking products over a wider range of carbon numbers according to a probability assigned to each bond. This modification leads to product distributions closer to the experimental data, particularly for iso-paraffins in the C_5 – C_{14} range, and therefore improves the prediction of the gasoline and kerosene cuts. Two additional thermodynamic refinements are introduced at the same time. First, the Langmuir adsorption constants K_L are made dependent on both carbon number and temperature through an exponential expression involving an adsorption enthalpy ΔH_{ads} , expressed as a linear function of chain length. Second, the previous carbon-number-only correlation for the isomerization equilibrium constant K_{eq} is replaced by a temperature-dependent formulation derived from a dedicated thermodynamic study of n-paraffin isomerization equilibria, while retaining its dependence on carbon number [58]. Together, the breakage probability function, the temperature-dependent adsorption constants, and the revised isomerization equilibrium constants improve both the predicted product distributions and the model response to changes in operating conditions, especially temperature.

The authors therefore show that VLE, realistic C–C bond breakage statistics, and temperature-dependent thermodynamic terms must be considered together to obtain reliable predictions across the investigated operating range. The remaining major simplification is the use of a simplified LHHW rate expression, which is addressed in the subsequent development of the model.

1.4.3.3.5 Gambaro et al. (2011): Complete LHHW rate expression [7]

Gambaro et al. (2011) closes the series by testing the complete form of the LHHW rate expression, whose simplified version, valid for paraffins up to approximately C_{16} where carbenium ion concentrations at acid sites are negligible [65], is no longer justified for the heavy FT wax feed. Two models are developed: Model 1 retains the assumptions of Gamba et al. [6] but replaces the exponential Langmuir constants with a hyperbolic tangent, whose sigmoidal shape reflects the transition from gas-phase to liquid-phase behaviour with increasing chain length, and uses power-law expressions for activation energies and pre-exponential factors (15 parameters). Model 2 adds a K_{PD} term to the denominator of the rate expressions, representing the product of the dehydrogenation and protonation equilibrium constants, which accounts for the non-negligible carbenium ion concentration at the active sites for long-chain paraffins (17 parameters).

Statistical model discrimination selects Model 2 as superior, with R^2 improving from 0.958 to 0.965; the improvement is confirmed not to result from the higher number of parameters alone. The complete LHHW rate form is shown to be necessary for long-chain paraffins, where the system behaviour differs fundamentally from the light paraffins on which the simplified form was originally validated. Predicted product distributions agree well with experimental data across a conversion range of 8–99%, particularly for kerosene (C_{10} – C_{14}), gas oil (C_{15} – C_{22}), and atmospheric residue (C_{22+}) fractions.

1.4.4 Operating conditions

The hydrocracking of Co-LTFT waxes is governed by four operating parameters: temperature, total pressure, hydrogen-to-wax mass ratio, and weight hourly space velocity (WHSV). Table 1.4 reports the ranges found in the literature for these four parameters. All four have a significant effect on C_{22+} conversion, with temperature and WHSV playing the dominant role, followed by pressure and H_2 /wax ratio [4, 45].

Table 1.4: Range of operating conditions for Co-LTFT wax hydrocracking reported in the literature [5, 17].

Parameter	Range	Unit
Temperature	305–375	°C
Pressure	35–80	bar
H_2 /wax	0.06–0.15	kg/kg
WHSV	1–4	$kg_{n-C,in} kg_{cat}^{-1} h^{-1}$

The operating window is bounded by physical and chemical constraints on each parameter. Temperature is the most influential variable: too low, conversion of heavy waxes remains insufficient; too high, consecutive cracking reactions become dominant and the middle distillate yield decreases in favour of lighter fractions [4, 45]. Conversion decreases with pressure: increasing total pressure raises the hydrogen fugacity in the ADSORB denominator, affecting negatively the dehydrogenation equilibrium and reducing the concentration of reactive olefinic intermediates at the acid sites [5, 6, 45]; the upper bound therefore reflects a trade-off between reactor operability and reaction rate. The H_2 /wax ratio promotes conversion through its effect on vapour-liquid equilibrium, higher ratios shift heavy hydrocarbons toward the vapour phase, increasing their fugacity, but this effect saturates beyond a certain value as the vapour fraction approaches its thermodynamic limit [4]. Finally, WHSV is inversely related to the contact time available for reaction. High WHSV values correspond to short contact times, which limit the conversion of the wax fraction, whereas low WHSV values correspond to longer contact times, increasing the exposure of intermediate products to secondary cracking and thereby reducing middle-distillate selectivity [45].

1.4.5 Reactor configurations

1.4.5.1 Trickle-bed as the industrial standard

Hydrocracking is a three-phase process: a gas phase composed of hydrogen and vaporised light hydrocarbons, a liquid phase composed of the heavy hydrocarbon feed, and a solid catalyst phase must be brought into contact simultaneously under elevated pressure and temperature.

In conventional petroleum hydrocracking, the trickle-bed reactor, a fixed bed of catalyst particles over which gas and liquid flow cocurrently downward, is the industrial standard. It is used in essentially all commercial units, the only notable exceptions being a few ebullated-bed processes (H-Oil, Hy-C) dedicated to very heavy residues that require continuous catalyst replacement [77].

For Co-LTFT wax hydrocracking, the same configuration prevails. Bouchy et al. report that the trickle bed is the reactor technology adopted across all forms of hydrocracking relevant to wax upgrading, from conventional petroleum hydrocracking to mild hydrocracking and FT wax hydrocracking alike [32]. The existing commercial FT wax hydrocrackers, operated at Sasol's Oryx GTL and Shell's Pearl GTL facilities, likewise rely on fixed-bed trickle-flow technology to convert Co-LTFT wax into middle distillates. No industrial-scale unit dedicated to Co-LTFT wax conversion within a Power-to-Kerosene (PtK) scheme has yet been built, since the upstream PtK chain itself remains at pilot stage; the trickle bed nevertheless constitutes the natural reference configuration for any future scale-up.

1.4.5.2 Operating principle

In a trickle-bed reactor, as illustrated in Figure 1.6, the hydrocarbon feed and hydrogen are introduced together at the top of the bed. The liquid hydrocarbon fraction percolates downward under gravity, forming a thin film that wets the external surface of the catalyst pellets and trickles from particle to particle along the packing. The gas phase flows cocurrently downward through the remaining void fraction, in continuous contact with the liquid film and the catalyst surface. Mass transfer thus proceeds through a sequence of resistances: hydrogen must dissolve from the gas into the liquid film, diffuse across this film to the catalyst surface, and then penetrate the porous network of the pellet to reach the active sites, where the dissolved hydrocarbons undergo isomerization and cracking.

Under typical hydrocracking conditions, the feed is itself partially vaporised: lighter hydrocarbons escape as vapour while the liquid phase, enriched in heavy C_{22+} paraffins, covers the catalyst particles and constitutes the primary reaction medium [32]. This vapour-liquid distribution, strongly dependent on temperature, pressure and H_2 /wax ratio as discussed in Section 1.4.4, is a defining characteristic of trickle-bed operation for this system.




Figure 1.6: Schematic representation of a trickle-bed reactor operated in co-current gas-liquid downflow mode.

1.4.5.3 Suitability for the Co-LTFT wax feed

The trickle bed is particularly well-suited to Co-LTFT wax. The clean nature of the feed, free of sulfur, nitrogen, and metal contaminants, eliminates the need for continuous catalyst replacement or regeneration, making a fixed-bed configuration appropriate; the long catalyst lifetime achievable under these conditions further supports this choice. All kinetic studies reviewed in Section 1.4.3.3 were accordingly conducted on bench-scale trickle beds operated in down-flow mode, loaded with finely powdered catalyst so as to suppress intraparticle and interphase diffusional resistances and approach ideal plug-flow behaviour [4].

Chapter 2

Hydrocracking reactor modelling

The state of the art reviewed in Chapter 1 identified hydrocracking over a bifunctional catalyst as the most appropriate technology for converting the heavy fraction of Co-LTFT syncrude into the kerosene-range hydrocarbons targeted by the Power-to-Kerosene process developed at the University of Liège. It further established the methodological framework on which any quantitative description of this conversion must rest: the LHHW kinetic formalism, the bifunctional reaction scheme reducing to hydroisomerization followed by cracking, and the Pellegrini-Gambaro modelling series as the most complete body of work available for the kinetic description of Co-LTFT wax hydrocracking. The present chapter translates this framework into a reactor model and applies it to the specific feedstock of the ULiège process.

The objective pursued here is twofold. The first is to construct, implement, and prepare for validation a kinetic model of the hydrocracking reactor that resolves the product stream at the component level, a resolution made necessary by the fact that the kerosene cut of interest, defined in this work as the C_8 – C_{16} fraction, cannot be evaluated from a lumped description without committing to a fixed cut definition. The second is to use this model as a predictive tool: to identify the operating conditions that maximise the kerosene yield for the ULiège feed, and to translate the resulting operating point into the physical dimensions of the reactor.

The chapter is organised in five sections, following the logical sequence of the modelling effort. Section 2.1 characterises the stream entering the reactor, starting from the Fischer-Tropsch effluent provided by the upstream ULiège process, and details the lumping strategy and the olefin-hydrogenation pre-treatment applied to it. Section 2.2 justifies the selection of the all-components model of Pellegrini et al. (2007) among the candidates of the Pellegrini-Gambaro series, and motivates the choice of Python over Aspen Plus as the implementation environment. Section 2.3 develops the kinetic model itself: the component-level material balance equations, the parameter-estimation procedure devised to recover the kinetic constants that the original publication does not report, and the thermodynamic assumption under which the partial-pressure form of the rate expressions remains admissible. Section 2.4 defines the operating window explored and sets out the structured search procedure used to locate the kerosene-maximising operating conditions. Finally, Section 2.5 establishes the methodology that converts these operating conditions into a reactor geometry consistent with the hydrodynamic criteria required for plug-flow behaviour.

It should be emphasised that the present chapter is concerned exclusively with the construction and description of the model. Its validation against experimental data, its application to the ULiège inlet stream, and the discussion of the results it produces are deferred to Chapter 3.

2.1 Feed stream characterisation

2.1.1 Origin and composition of the feed

The ULiège process into which the present work integrates is shown in Figure 2.1, in simplified form and restricted to the units relevant here. The upstream section, composed of the reverse water-gas shift (RWGS) reactor, the low-temperature cobalt-based Fischer-Tropsch (LTFT) reactor and the associated condensation and flash-separation steps, was developed by A. Rouxhet and A. Morales and defines the stream entering this study.




Figure 2.1: Simplified block flow diagram of the relevant section of the ULiège process and the hydrocracking unit modelled in this work.

The inlet stream considered is that provided by doctoral researcher A. Rouxhet, taken at the outlet of the flash separator operating at 150 °C and 20 bar, whose purpose is to maximise the condensation of the Fischer-Tropsch products heavier than C_{16} . Between the flash separator and the hydrocracker, the stream was assumed to be routed through a pump followed by a heat exchanger, bringing the feed to the hydrocracking operating conditions. Given the limited time available to carry out this work, neither the energy integration of the process nor the detailed modelling of these two units could be undertaken. These aspects therefore remain important perspectives for future work. Finally, to make the conversion sequence easier to follow, the olefin hydrogenation step is represented as a separate stoichiometric reactor in Figure 2.1, although in practice it occurs within the hydrocracker, as explained in Section 2.1.3.

As detailed in Table 2.1, the stream exiting the flash separator is composed of water, CO_2 , CO , hydrogen, n-paraffins, and olefins. The detailed thermodynamic state and molar composition per species are given in Appendix A.2. In the context of hydrocracking reactor modelling, only paraffins, olefins, and hydrogen act as reactive species. It is therefore assumed that water, CO_2 , and CO are removed upstream of the hydrocracking reactor. Accordingly, all molar fractions reported in this chapter are expressed relative to the FT hydrocarbon product stream only.

Table 2.1: Stream properties at the flash separator outlet before and after removal of non-reactive species.

Property	Unit	Raw stream	FT hydrocarbon product stream
Operating conditions			
Temperature	°C	150	150
Pressure	bar	20	20
Phase	—	Liquid	Liquid
Total molar flow	kmol h ⁻¹	1.0×10^{-3}	4.83×10^{-4}
Molar composition			
H ₂ O	mol%	50.171	—
CO ₂	mol%	0.888	—
CO	mol%	0.283	—
H ₂	mol%	0.353	—
<i>n</i> -Paraffins	mol%	46.522	96.308
Olefins	mol%	1.784	3.692

It should also be noted that, compared to the typical syncrude compositions reported in the literature for LTFT processes (see Appendix A.1), the stream considered in this work contains neither oxygenated compounds nor branched species. This results from a simplifying assumption introduced by A. Rouxhet, justified by the negligible amounts of such species under the conditions considered; the reader is referred to his thesis, the publication of which is expected later this year, for further details.

2.1.2 Lumping strategy

The kinetic model of Pellegrini et al. (2007), on which this work is based (see Section 2.2.2), is formulated for a feedstock of *n*-paraffins ranging from C₁ to C₇₀. The ULiège stream, as originally provided, extended only up to a C₃₀₊ lump (see Table 2.1); to match the carbon-number range of this model, A. Rouxhet extended this composition into a full C₁-C₇₀₊ distribution. All compounds heavier than C₆₉ are then grouped into a single lump designated C₇₀₊. As this lump represents only 0.012 mol% of the stream, this simplification has a negligible influence on the results. The paraffin and olefin compositions of the inlet hydrocarbon stream are shown in Figures 2.2a and 2.2b, respectively.




Figure 2.2: Initial molar-fraction distribution of the ULiège hydrocarbon inlet stream.

2.1.3 Olefin hydrogenation pre-treatment

As shown in Figure 2.2b, the olefin content of the inlet stream is not negligible. Owing to their considerably higher reactivity towards hydrogen than *n*-paraffins, discussed in Section 1.4.1, the olefins are supposed instantaneously hydrogenated upon entering the hydrocracking reactor. The heat released by

this exothermic saturation reaction partially compensates the endothermic cracking reactions, contributing to the near-isothermal behaviour of the process, as reported in Section 1.3.4.

Because the olefins are supposed to be converted to n-paraffins immediately, hydrocracking models are predominantly developed for feeds composed exclusively of n-paraffins: either a dedicated upstream hydrogenation unit has already removed the olefins [78,79], or their instantaneous conversion within the reactor is assumed directly [5].

In the present work, the olefin hydrogenation is represented as a distinct stoichiometric reactor placed upstream of the hydrocracker [78,79], so that the conversion of the olefins into n-paraffins appears explicitly in the process description. The olefins are taken to be fully converted to the corresponding n-paraffins according to Reaction (2.1), a conversion of 100% being adopted in line with their high hydrogenation reactivity.



Because this conversion is total and instantaneous, representing it as a separate reactor is numerically equivalent to assuming the olefins to be hydrogenated at the hydrocracking reactor inlet; the separate block is therefore a choice of representation rather than a distinct modelling step. In an industrial setting, a dedicated hydrogenation unit would not be justified, the reaction taking place directly within the hydrocracking reactor.

Once the olefins have been fully hydrogenated, the stream consists entirely of n-paraffins and can undergo hydrocracking in the presence of hydrogen. This n-paraffinic stream constitutes the inlet to the hydrocracking reactor for all simulations presented in Chapter 3, where its detailed composition is reported.

2.2 Model selection and implementation framework

2.2.1 Choice of Python as modelling tool

The upstream Power-to-Kerosene process, comprising the RWGS reactor, the Fischer-Tropsch reactor, and the associated separation units, was previously modelled in Aspen Plus [80] and Aspen Custom Modeler [81] by the Chemical Engineering team at ULiège, with contributions from doctoral researchers A. Rouxhet and A. Morales [82–86]. Aspen is a commercial process simulation software package providing thermodynamic property methods and unit-operation models. In the ULiège's model, the hydrocarbon fraction is represented using standard Aspen Plus components, with heavier paraffins aggregated into a single n-C₃₀₊ lump and olefins into a C₁₀₊ lump, as reported in the Appendix A.2. For the purposes of hydrocracking modelling, however, this level of aggregation is judged insufficiently detailed: a faithful representation of the cracking kinetics requires individual species tracking up further (such as up to C₇₀).

Extending the Aspen Plus model to this resolution raises two obstacles. First, the Aspen Plus thermodynamic database provides property data only up to n-C₅₀ for n-paraffins and iso-C₂₁ for iso-paraffins; heavier species must be defined as user-specified pseudo-components whose properties rely on correlation-based estimates of reduced accuracy. Second, the LHHW kinetic framework with 138 components is not well-suited to the Aspen Plus reaction block interface: each of the 138 rate equations must be entered individually, five kinetic parameters must be specified per component, and the resulting adsorption matrix is of a size that exceeds the practical limits of the solver. For these reasons, Python was selected as the modelling environment for the hydrocracking reactor.

Python offers full control over the governing equations, allowing the material balances, kinetic rate expressions, and numerical integration to be implemented directly, without being constrained by predefined block structures, fixed kinetic models, or the internal solver logic. It also provides greater flexibility and computational speed for the parametric studies presented in Chapter 3.

However, the primary limitation of Python relative to Aspen Plus is the absence of a thermodynamic library of comparable scope. For a mixture of n-paraffins and iso-paraffins spanning C₁ to C₇₀, all thermodynamic properties and phase equilibrium calculations must be implemented manually. A vapour-liquid equilibrium framework based on the Soave-Redlich-Kwong equation of state was investigated over the

course of several weeks, using Wilson K-value correlations to initialise the Rachford-Rice flash algorithm and the SRKMIX library for fugacity coefficient evaluation. Despite sustained effort, this approach did not yield satisfactory results: the combination of numerical instabilities in the thermodynamic library with the large ADSORB denominator of the LHHW rate expression led to a severe underestimation of cracking conversion. Although a fully rigorous fugacity-based implementation on Python may well be achievable, the complexity of the problem exceeded what could be addressed within the time frame of this thesis. A partial-pressure model was therefore retained, whose validity conditions are discussed in Section 2.3.3.

It should be noted that Aspen Custom Modeler (ACM) would have represented a technically appropriate alternative, offering the flexibility to implement custom equation-based models and thereby removing the modelling-interface limitation of Aspen Plus [81]. The thermodynamic database limitation would nonetheless have persisted, since ACM shares the same property database as Aspen Plus. However, acquiring the necessary proficiency in ACM within the time constraints of the present work was not feasible, and Python was therefore preferred.

2.2.2 Selection of the all-components model of Pellegrini et al. (2007)

The kinetic model adopted in this work is based on the all-components model developed by Pellegrini et al. (2007) [3]. This choice was motivated by several considerations specific to the objectives and constraints of the present work.

First, the Pellegrini-Gambaro series is built on the LHHW formalism (Section 1.4.3.1), which derives the hydroisomerization and cracking rates from the bifunctional reaction pathway and from competitive surface adsorption. This physical basis is what makes the series the reference kinetic description for Co-LTFT wax hydrocracking, and the reason it is retained here.

Second, among the Pellegrini-Gambaro series, the 2007 all-components model is the most advanced formulation that does not require a fugacity-based vapour-liquid equilibrium treatment. The later models (the VLE formulation of 2008 [4], the breakage-probability development [6], and the complete LHHW rate expression [7]) all rely on fugacities, whose reliable computation could not be achieved within the modelling environment retained (Section 2.2.1). The 2007 model is therefore the most complete model of the series compatible with the partial-pressure approach adopted in this work, while still reproducing the experimental data with sufficient accuracy for the present objective (Chapter 3).

Third, since the ultimate objective is to maximise the production of the C₈-C₁₆ kerosene cut, a model describing the evolution of each individual component is required. The definition of the kerosene fraction is not uniform across the literature, and lumped models built around a fixed cut definition cannot be straightforwardly adapted to the C₈-C₁₆ boundary adopted in this work. Component-level material balances, by contrast, allow any cut to be evaluated in post-processing without any modification to the model structure. Lumped models are therefore excluded on this basis.

This model nonetheless presents two structural limitations: the single vapour-phase partial-pressure assumption and the midpoint cracking stoichiometry. They are recognised by the original authors, who addressed them in their subsequent publications, and discussed in the present work in Chapter 3.

Description of the Pellegrini et al. (2007) model

The all-components model is built on the simplified kinetic scheme of Reaction (2.2), derived from the full Steijns-Froment network [71] by merging all branched isomers of carbon number i into a single lumped pseudo-component iso-C _{i} , as discussed in Section 1.4.3.3.2. Direct cracking of n-paraffins is neglected (see Section 1.4.1.1): the linear paraffin must first isomerize before cracking becomes accessible.



The reactor is modelled as an isothermal plug-flow reactor, described by Equation (2.3), where Y_i denotes the molar amount of species i per unit mass of inlet n-paraffins, τ the contact time (kg_{cat}·h/kg_{n-C_{in}}) and r_i is the net reaction rate of species i , expressed in kmol _{i} /kg_{cat}/h. Inter- and intra-phase diffusion

resistances are neglected consistent with the use of a powdered catalyst in the experimental campaign on which the model is calibrated.

$$\frac{dY_i}{d\tau} = r_i \quad (2.3)$$

The LHHW rate expressions for hydroisomerization and cracking are given by Equations (2.4) and (2.5) respectively, with Arrhenius temperature dependence for the kinetic constants $k_{\text{isom}}(i)$ and $k_{\text{cr}}(i)$.

$$r_{\text{isom}}(i) = \frac{k_{\text{isom}}(i) e^{-E_{\text{isom}}(i)/RT} \left(p_n(i) - \frac{p_{\text{iso}}(i)}{K_{\text{eq}}(i)} \right)}{\text{ADSORB}} \quad (2.4)$$

$$r_{\text{cr}}(i) = \frac{k_{\text{cr}}(i) e^{-E_{\text{cr}}(i)/RT} p_{\text{iso}}(i)}{\text{ADSORB}} \quad (2.5)$$

The competitive adsorption denominator ADSORB captures the inhibition of all surface reactions by competitive adsorption of all species on the catalyst, as well as the experimentally observed inhibition at high hydrogen pressure:

$$\text{ADSORB} = p_{\text{H}_2} \left[1 + \sum_{i=1}^{70} K_{L,n}(i) p_n(i) + \sum_{i=4}^{70} K_{L,\text{iso}}(i) p_{\text{iso}}(i) \right] \quad (2.6)$$

The model is presented in two parts. In the first, each kinetic constant takes a single fixed value per lump. The second part is an evolution of the first: each kinetic constant is expressed as a continuous function of the carbon number, obtained by fitting the discrete lumped values with simple functional forms (power-law for the frequency factors, logarithmic for the activation energies, exponential for the Langmuir constants). These continuous expressions remove the discontinuities at the lump boundaries and provide a more physically consistent description of the carbon-number dependence. Although the lump-constant variant already yields satisfactory agreement with the experimental data, the continuous-function variant constitutes the more complete formulation and is therefore the one retained in this work.

Two practical challenges arise in its implementation. First, the authors do not report the inlet composition used to obtain the experimental and model results, which complicates verification and reproduction. This difficulty can be addressed using the first variant: its lumped kinetic constants are reported explicitly for every lump, so the inlet composition is the only remaining unknown and can be recovered by requiring the lumped model to reproduce the reported outlet. This identification is carried out in Section 3.1.1. Second, the kinetic expressions used for the final calibration of the continuous-function variant are not given: the authors report only the final graphical results and the initial expressions used as starting points for the non-linear regression that yields the optimised parameters. The functional expressions and the methodology used to recover these final kinetic expressions are presented in Section 2.3.1.



2.3 Kinetic model implementation

2.3.1 Material balance equations


Based on the model described in Section 2.2.2, the full system of component material balances is implemented in Python. The 138 species, consisting of 70 n-paraffins, 67 iso-paraffins, and hydrogen, are tracked individually through a system of ordinary differential equations of the form of Equation (2.3), integrated numerically using the adaptive Runge-Kutta RK45 scheme. The right-hand sides r_i of Equation (2.3) are developed explicitly below for each species class, following the midpoint cracking stoichiometry and the LHHW rate expressions of Equations (2.4)-(2.6).

Stoichiometry

As discussed in Section 1.4.3.3.2, the cracking stoichiometry adopted in the model follows the midpoint approximation, in which each iso-paraffin of carbon number n undergoes primary cracking at the central C-C bond, yielding two paraffinic fragments. For an even carbon number, the two fragments are identical; for an odd carbon number, the two neighbouring fragments are produced. The global stoichiometric equations, which consume one mole of hydrogen per cracking event, are given by Reactions (2.7) and (2.8) for $n \geq 6$.



The treatment of light species constitutes an explicit modelling choice of Pellegrini et al. (2007) [3]. n -C₄ is allowed to isomerize to iso-C₄, while iso-C₄ does not crack. C₅ is the lightest iso-paraffin that undergoes cracking, but with a product distribution that deviates from the midpoint approximation: rather than producing the expected C₂ + C₃ pair, iso-C₅ cracks according to Reaction (2.9).



Carbon conservation and mole balance are verified: $0.5 \times (1 + 2 + 3 + 4) = 5$ carbon atoms and $0.5 \times 4 = 2$ fragments per cracking event. The species C₁, C₂, and C₃ are terminal: they undergo neither isomerization nor cracking and appear solely as products of heavier species.

System of ODEs

The material balances for all 138 species are given by Equations (2.10)-(2.18), following the formulation of Pellegrini et al. (2007) [3]. For n -paraffins C₄ to C₇₀, the balance reflects consumption by hydroisomerization only, since direct cracking of n -paraffins is neglected [3]:

$$\frac{dY_{nC_i}}{d\tau} = -\frac{k_{\text{isom}}(i) e^{-E_{\text{isom}}(i)/RT}}{\text{ADSORB}} \left(p_n(i) - \frac{p_{\text{iso}}(i)}{K_{\text{eq}}(i)} \right), \quad i = 4, \dots, 70 \quad (2.10)$$

As stated in the stoichiometry model, the production terms for C₁, C₂, and C₃ described by Equations (2.11)-(2.13), arise exclusively from the cracking of iso-C₅ and from the midpoint cracking of heavier species.

$$\begin{aligned} \frac{dY_{nC_i}}{d\tau} = & \frac{2k_{\text{cr}}(2i)e^{-E_{\text{cr}}(2i)/RT} p_{\text{iso}}(2i) + 0.5k_{\text{cr}}(2i-1)e^{-E_{\text{cr}}(2i-1)/RT} p_{\text{iso}}(2i-1)}{\text{ADSORB}} \\ & + \frac{k_{\text{cr}}(2i+1)e^{-E_{\text{cr}}(2i+1)/RT} p_{\text{iso}}(2i+1)}{\text{ADSORB}}, \quad i = 3 \end{aligned} \quad (2.11)$$

$$\frac{dY_{nC_i}}{d\tau} = \frac{0.5 k_{\text{cr}}(2i+1) e^{-E_{\text{cr}}(2i+1)/RT} p_{\text{iso}}(2i+1)}{\text{ADSORB}}, \quad i = 2 \quad (2.12)$$

$$\frac{dY_{nC_i}}{d\tau} = \frac{0.5 k_{\text{cr}}(4i+1) e^{-E_{\text{cr}}(4i+1)/RT} p_{\text{iso}}(4i+1)}{\text{ADSORB}}, \quad i = 1 \quad (2.13)$$

For iso-paraffins, the balance accounts for production by hydroisomerization, production by cracking of heavier iso-paraffins via the midpoint stoichiometry, and consumption by cracking:

$$\frac{dY_{\text{iso}C_i}}{d\tau} = -\frac{dY_{nC_i}}{d\tau} - \frac{k_{\text{cr}}(i) e^{-E_{\text{cr}}(i)/RT}}{\text{ADSORB}} p_{\text{iso}}(i), \quad i = 36, \dots, 70 \quad (2.14)$$

$$\begin{aligned} \frac{dY_{\text{iso}C_i}}{d\tau} = & -\frac{dY_{nC_i}}{d\tau} + \frac{2k_{\text{cr}}(2i) e^{-E_{\text{cr}}(2i)/RT} p_{\text{iso}}(2i) + k_{\text{cr}}(2i-1) e^{-E_{\text{cr}}(2i-1)/RT} p_{\text{iso}}(2i-1)}{\text{ADSORB}} \\ & - \frac{k_{\text{cr}}(i) e^{-E_{\text{cr}}(i)/RT} p_{\text{iso}}(i)}{\text{ADSORB}}, \quad i = 35 \end{aligned} \quad (2.15)$$

$$\begin{aligned} \frac{dY_{\text{iso}C_i}}{d\tau} = & -\frac{dY_{nC_i}}{d\tau} + \frac{2k_{\text{cr}}(2i) e^{-E_{\text{cr}}(2i)/RT} p_{\text{iso}}(2i) + k_{\text{cr}}(2i-1) e^{-E_{\text{cr}}(2i-1)/RT} p_{\text{iso}}(2i-1)}{\text{ADSORB}} \\ & + \frac{k_{\text{cr}}(2i+1) e^{-E_{\text{cr}}(2i+1)/RT} p_{\text{iso}}(2i+1) - k_{\text{cr}}(i) e^{-E_{\text{cr}}(i)/RT} p_{\text{iso}}(i)}{\text{ADSORB}}, \quad i = 5, \dots, 34 \end{aligned} \quad (2.16)$$

$$\begin{aligned} \frac{dY_{\text{iso}C_i}}{d\tau} = & -\frac{dY_{nC_i}}{d\tau} + \frac{2k_{\text{cr}}(2i) e^{-E_{\text{cr}}(2i)/RT} p_{\text{iso}}(2i) + k_{\text{cr}}(2i-1) e^{-E_{\text{cr}}(2i-1)/RT} p_{\text{iso}}(2i-1)}{\text{ADSORB}} \\ & + \frac{k_{\text{cr}}(2i+1) e^{-E_{\text{cr}}(2i+1)/RT} p_{\text{iso}}(2i+1) + 0.5 k_{\text{cr}}(i+1) e^{-E_{\text{cr}}(i+1)/RT} p_{\text{iso}}(i+1)}{\text{ADSORB}}, \quad i = 4 \end{aligned} \quad (2.17)$$

Finally, the hydrogen balance follows directly from the cracking stoichiometry, since each cracking event consumes exactly one mole of hydrogen. As iso-C₄ does not crack, the summation starts from $i = 5$:

$$\frac{dY_{\text{H}_2}}{d\tau} = -\sum_{i=5}^{70} \frac{k_{\text{cr}}(i) e^{-E_{\text{cr}}(i)/RT}}{\text{ADSORB}} p_{\text{iso}}(i) \quad (2.18)$$

The competitive adsorption term ADSORB is defined as:

$$\text{ADSORB} = p_{\text{H}_2} \left[1 + \sum_{i=1}^{70} K_{L,n}(i) p_n(i) + \sum_{i=4}^{70} K_{L,\text{iso}}(i) p_{\text{iso}}(i) \right] \quad (2.19)$$

The partial pressure of each species is computed from the molar amounts Y_i and the total pressure P as:

$$p_i = \frac{Y_i}{\sum_j Y_j} \cdot P \quad (2.20)$$

where the summation runs over all species, including hydrogen. All variables and parameters involved in this system are defined in Table 2.2.

Table 2.2: Definition of variables and parameters appearing in the material balance equations.

Symbol	Definition	Unit
Y_i	Molar amount of species i per unit mass of inlet n-paraffins	$\text{kmol kg}_{n\text{-C},\text{in}}^{-1}$
τ	Contact time	$\text{kg}_{\text{cat}} \text{ h kg}_{n\text{-C},\text{in}}^{-1}$
$p_n(i)$	Partial pressure of n-paraffin C_i	Pa
$p_{\text{iso}}(i)$	Partial pressure of iso-paraffin C_i	Pa
p_{H_2}	Partial pressure of hydrogen	Pa
P	Total pressure	Pa
$k_{\text{isom}}(i)$	Pre-exponential factor for isomerization of C_i	$\text{kmol h}^{-1} \text{ kg}_{\text{cat}}^{-1}$
$k_{\text{cr}}(i)$	Pre-exponential factor for cracking of C_i	$\text{kmol h}^{-1} \text{ kg}_{\text{cat}}^{-1}$
$E_{\text{isom}}(i)$	Activation energy for isomerization of C_i	J mol^{-1}
$E_{\text{cr}}(i)$	Activation energy for cracking of C_i	J mol^{-1}
$K_{\text{eq}}(i)$	Lumped isomerization equilibrium constant for C_i	-
$K_{L,n}(i)$	Langmuir adsorption constant for n-paraffin C_i	Pa^{-1}
$K_{L,\text{iso}}(i)$	Langmuir adsorption constant for iso-paraffin C_i	Pa^{-1}
R	Universal gas constant	$\text{J mol}^{-1} \text{ K}^{-1}$
T	Reactor temperature	K

Kinetic constant expressions

As explained in Section 2.2.2, the second part of the Pellegrini et al. (2007) model expresses each kinetic constant as a continuous function of carbon number n_c . To understand how these expressions are established, it is necessary to retrace the procedure followed by the authors. In the lumped model, each kinetic parameter is determined by non-linear regression against the experimental dataset, yielding a single fixed value per lump. These lump values, reported in Appendix A.3, provide four to five reference points per parameter. In the second part of the article, these discrete lump values serve as calibration points to fit continuous functions of n_c by non-linear regression. Since only four or five data points are available per parameter, the functional forms are simple: power laws for the frequency factors, logarithmic expressions for the activation energies, and exponential or polynomial expressions for the adsorption and equilibrium constants. The resulting expressions are given in Equations (2.21)-(2.27).

$$K_{L,n}(n_c) = 0.1 \cdot \exp(0.39 n_c), \quad n_c = 1, \dots, 70 \quad (2.21)$$

$$K_{L,\text{iso}}(n_c) = 0.1 \cdot \exp(0.38 n_c), \quad n_c \geq 4 \quad (2.22)$$

$$k_{\text{isom}}(n_c) = 5 \times 10^{16} \cdot n_c^{7.9241}, \quad n_c \geq 4 \quad (2.23)$$

$$k_{\text{cr}}(n_c) = 2 \times 10^{18} \cdot n_c^{5.7677}, \quad n_c \geq 5 \quad (2.24)$$

$$E_{\text{isom}}(n_c) = 22253 \ln(n_c) + 118350, \quad n_c \geq 4 \quad (2.25)$$

$$E_{\text{cr}}(n_c) = 29806 \ln(n_c) + 86006, \quad n_c \geq 5 \quad (2.26)$$

$$K_{\text{eq}}(n_c) = 1.1268 n_c^2 - 3.1914 n_c + 3589.1, \quad n_c \geq 4 \quad (2.27)$$

These expressions, however, do not represent the final kinetic parameters used in the model. The authors use them solely as initial values for a second non-linear regression, in which the reactor is simulated at each iteration and the expressions are adjusted by least-squares minimisation, optimised directly against the full experimental dataset. This second regression adjusts the coefficients within each functional form, for instance, the prefactor and exponent in the power-law expressions, or the slope and intercept in the

logarithmic ones, to minimise the deviation between model predictions and experimental measurements. It is the result of this second regression that constitutes the truly ameliorated model, and that Pellegrini et al. (2007) use to illustrate the final product distributions in their article. Unfortunately, these final optimised coefficients are not reported in the publication, which represents one of the two main practical challenges in reproducing this model. The methodology developed in this work to recover this consistent set of kinetic parameters is described in the following section, and the resulting values are reported in Chapter 3.

2.3.2 Kinetic parameter estimation

As established in the previous section, the expressions given in Equations (2.21)-(2.27) constitute the starting point of Pellegrini's non-linear regression, not the final optimised parameter values. Since these final expressions are not reported in the publication, a dedicated parameter estimation procedure is developed in this work to recover a consistent set of kinetic parameters. Specifically, the objective is to identify the 15 coefficients $\theta = (A, B, D, E, F, G, H, I, J, L, M, N, O, P, Q)$ appearing in the functional forms:

$$K_{L,n}(n_c) = A \cdot \exp(B \cdot n_c) \quad (2.28)$$

$$K_{L,iso}(n_c) = D \cdot \exp(E \cdot n_c) \quad (2.29)$$

$$K_{eq}(n_c) = F n_c^2 - G n_c + H \quad (2.30)$$

$$k_{isom}(n_c) = I \cdot n_c^J \quad (2.31)$$

$$k_{cr}(n_c) = L \cdot n_c^M \quad (2.32)$$

$$E_{isom}(n_c) = N \ln(n_c) + O \quad (2.33)$$

$$E_{cr}(n_c) = P \ln(n_c) + Q \quad (2.34)$$

These parameters are regressed using the outlet product distributions reported by Pellegrini et al. (2007) [3]. Although Pellegrini published both the optimised model output curves and the experimental data points in the same figures, making it in principle possible to target either, this work deliberately optimises against the model output curves of Figures 2.3a and 2.3b. The reason is that the primary objective is to reproduce Pellegrini's published model as faithfully as possible, in order to establish a reliable and validated foundation before applying it to the ULiège inlet stream. Targeting the experimental data directly would conceptually replicate Pellegrini's approach, but would not guarantee convergence to the same solution, the optimisation could settle on a different parameter set that fits the data equally well but diverges from Pellegrini's model. By contrast, targeting the model curves ensures that the reproduced parameter set is anchored to a known, published reference. The article only provides the final distribution with only one set of operating conditions. Overfitting is therefore a known risk: a parameter set that fits the target curves perfectly but has no physical grounding may produce unreliable predictions under different operating conditions. The regularisation strategy described below is introduced precisely to address this concern. For the same reason, it cannot be excluded that Pellegrini's model is itself overfitted; this limitation is intrinsic to the available data and lies beyond the scope of the regularisation strategy.

Target curves

The target distributions are obtained by digitising Figures 2.3a and 2.3b of Pellegrini et al. (2007) [3], which report the outlet molar fractions of n-paraffins and iso-paraffins under conditions $T = 367 \text{ }^\circ\text{C}$, $P = 53.75 \text{ bar}$, $H_2/\text{wax} = 0.1275 \text{ kg/kg}$, $WHSV = 2.5 \text{ kg}_{n-C, \text{in}} \text{ kg}_{\text{cat}}^{-1} \text{ h}^{-1}$. These two "model with optimised K" curves together constitute the reference against which the optimised parameters are evaluated.




Figure 2.3: Final distribution obtained by Pellegrini et al. under conditions $T = 367$ °C, $P = 53.75$ bar, $H_2/wax = 0.1275$ kg/kg, $WHSV = 2.5$ $kg_{n-C,in} kg_{cat}^{-1} h^{-1}$ using the final kinetic constants expressed as functions of carbon number [3].

Classical objective function and ill-posedness

A natural first approach is to minimise the sum of squared errors (SSE) between the simulated and target distributions [87]:

$$\mathcal{L}_{SSE}(\theta) = \sum_{i=1}^{70} \left(x_{n,i}^{sim} - x_{n,i}^{target} \right)^2 + \sum_{i=4}^{70} \left(x_{iso,i}^{sim} - x_{iso,i}^{target} \right)^2 \quad (2.35)$$

where θ denotes the vector of 15 parameters to be optimised. However, this problem is ill-posed: with 15 free parameters and only two output curves to match, the problem is over-parameterised and multiple very different parameter sets yield virtually identical SSE values. This non-uniqueness makes the classical SSE minimisation unreliable, as the optimiser may converge to physically meaningless parameter combinations.

Tikhonov regularisation

To address the identifiability problem, a Tikhonov regularisation term is added to the objective function [88]. This penalty term penalises relative deviations of the optimised parameters from the reference values of Pellegrini [3], thereby constraining the solution to remain physically consistent with the original model:

$$\mathcal{L}(\theta) = \mathcal{L}_{SSE}(\theta) + \lambda \sum_{j \in \mathcal{R}} w_j \left(\frac{\theta_j - \theta_j^{ref}}{\theta_j^{ref}} \right)^2 \quad (2.36)$$

where λ is the global regularisation strength, setting the overall trade-off between reproducing the reference distributions and keeping the parameters close to their reference values, and w_j are parameter-specific weights that distribute this overall strength unequally across the regularised subset $\mathcal{R}$, assigning a stronger relative penalty to the parameters according to their physical importance and their observed tendency to deviate; and θ_j^{ref} are the reference values taken from Pellegrini's starting expressions. The weighting is set from observation: during fitting, the optimiser tends to move certain parameters strongly away from their reference values, so higher weights are assigned to the parameters of K_{eq} , E_{cr} and the cracking rate exponent, which are the most sensitive to over-fitting, while $K_{L,n}$ and $K_{L,iso}$ are left unconstrained, as they showed no significant deviation during fitting. The choice of retaining the functional forms of Pellegrini's expressions is supported by the observation that Pellegrini et al. (2008) [4], in extending the model to include vapour–liquid equilibrium, preserve the same functional forms while adjusting only the coefficients.

Optimisation algorithm

To avoid getting trapped in local minima, the optimisation is carried out in three successive stages of increasing precision. The pre-exponential factors $K_{L,n}$, $K_{L,iso}$, k_{iso} , and k_{cr} span several orders of magnitude and are therefore optimised in logarithmic space, while the remaining parameters are handled in their natural space. The three stages are:

1. A coarse Nelder-Mead search [89] using a fast fixed-step RK4 integrator with 60 steps, allowing rapid exploration of the parameter space;
2. A Powell refinement [90] with 80 integration steps, converging toward a more precise minimum;
3. A final Powell refinement with 150 integration steps, ensuring convergence to a tight tolerance.

The result of this procedure is a set of 15 optimised coefficients whose values are reported in Chapter 3.

Extension to experimental data

The same regularised optimisation framework is subsequently applied using the experimental data of Pellegrini et al. (2007) as the target, rather than the model curves, conceptually replicating Pellegrini's approach. This yields a second set of kinetic parameters directly regressed against experiment, enabling a comparison between the two approaches. The corresponding results are also reported in Chapter 3.

2.3.3 Thermodynamic assumption

As established in Section 2.2.1, the vapour-liquid equilibrium framework required to compute fugacities could not be implemented reliably in Python. The model therefore operates with partial pressures in place of fugacities throughout the LHHW rate expressions of Equations (2.4)-(2.6). This substitution is physically justified when the reactive mixture remains overwhelmingly in the vapour phase, so that deviations between partial pressures and fugacities remain negligible. Although this approximation is reasonable for lighter hydrocarbons, deviations may occur for heavier compounds, which tend to partition preferentially into the liquid phase, as illustrated by the component-level recovery ratios in the vapour phase reported by Pellegrini et al. (2007) [3] in their Figure 2.4. Therefore, the accuracy of the model may decrease under conditions where liquid-phase effects become significant.




Figure 2.4: Recovery ratio in vapour phase for each component, plotted by Pellegrini et al. [3].

The validity condition associated with this assumption concerns the overall vapour mole fraction of the reactive mixture. Pellegrini et al. (2007) justify the use of partial pressures by showing that the vapour phase is predominant on a molar basis throughout the reactor under their operating conditions and feed composition. Hosukoglu et al. [72], applying the same kinetic model, further state that the vapour fraction does not fall below 96% in typical hydrocracking conditions, attributing this to the high hydrogen mole fraction in the feed ($\approx 95\%$), which remains in the vapour phase and maintains the overall mixture vapour

fraction at high levels. On the basis of these two observations, a minimum overall vapour mole fraction of 96% is adopted in this work as the validity criterion for the partial-pressure assumption. This condition must therefore be verified for every set of operating conditions considered, and any condition yielding a vapour fraction below this threshold is excluded from the analysis, as the partial-pressure assumption would no longer hold.

Verification using Aspen Plus

The verification is carried out using Aspen Plus with the Soave-Redlich-Kwong (SRK) equation of state [4], which provides a thermodynamic framework of sufficient scope to evaluate the vapour-liquid equilibrium of the mixture under the operating conditions of interest. Since the Aspen Plus property database does not include n-paraffins beyond n-C₅₀ as standard components, the heavier species (C₅₁ to C₇₀) are defined as user-specified pseudo-components. Two properties are required to define each pseudo-component: the molar mass and the normal boiling point.

The molar mass of each n-paraffin C_{n_c} is calculated from its molecular formula C_{n_c}H_{2n_c+2}:

$$M(n_c) = 12.011 n_c + 1.008 (2n_c + 2) \quad [\text{g/mol}] \quad (2.37)$$

The normal boiling point is estimated by a quadratic regression established on the known boiling points of n-paraffins, extrapolated to heavier species:

$$T_b(n_c) = -0.11729 n_c^2 + 16.2024 n_c + 71.2989 \quad [^\circ\text{C}] \quad (2.38)$$

The vapour fraction of the inlet mixture at operating conditions specified can therefore be evaluated. The verification results are reported in Chapter 3.

2.4 Operating conditions and optimal set identification

The kinetic model validated in Chapter 3 provides, for any combination of the four operating parameters, the outlet carbon-number distribution of the hydrocracking reactor. It can therefore be used to search for the optimal operating conditions that maximise the kerosene yield. This section describes the methodology developed for that purpose. Several modelling choices remain open at this stage: which fraction of the ULiège stream is sent to the reactor, how far the model is extrapolated beyond its calibration range, and how strictly the notion of optimal is defined. The procedure is therefore applied to a structured set of cases rather than to a single configuration. These choices are introduced in Section 2.4.1; the search procedure common to every case is described in Section 2.4.2; and the corresponding results are reported and discussed in Chapter 3.

2.4.1 Operating window and explored cases

Operating parameters and exploration domains

The four operating parameters governing the hydrocracking reactor are the temperature T , the total pressure P , the weight hourly space velocity WHSV, and the hydrogen-to-wax mass ratio H₂/wax. Two exploration domains are considered in this work, as summarised in Table 2.3.

Table 2.3: Operating domains explored in this work.

Domain	T (°C)	P (bar)	WHSV (kg _{n-C,in} kg _{cat} ⁻¹ h ⁻¹)	H ₂ /wax (kg/kg)
Pellegrini calibration	343-375	35-60	1-3	0.06-0.15
Extended physical	305-375	35-80	1-4	0.06-0.15

The first domain corresponds to the experimental conditions on which Pellegrini et al. (2007) [3] estimated the kinetic parameters used in this work. Operating points falling within this domain therefore lie within the calibration range of the model, and the corresponding predictions can be expected to carry a good degree of reliability. The second domain extends to the broader range of operating conditions reported in the literature for Co-LTFT wax hydrocracking (Section 1.4.4); exploring it inevitably entails extrapolating the model beyond its calibration range, and the corresponding results must therefore be interpreted with caution. Every case described below is solved independently on both domains, so that the cost of the extrapolation can be assessed by direct comparison.

Feed streams considered

The search is also carried out for two distinct feed streams, corresponding to two different ways of integrating the hydrocracking unit into the ULiège Power-to-Kerosene process.

The first is the complete hydrocarbon ULiège stream characterised in Section 2.1, spanning C_1 to C_{70+} after olefin hydrogenation. This stream already contains a substantial kerosene fraction of 33.8 wt-% of C_8 - C_{16} , inherited from the upstream Fischer–Tropsch reactor. The quantity maximised when this stream is processed is the C_8 - C_{16} mass fraction at the reactor outlet $w_{C_8-C_{16}}^{\text{out}}$. Since the inlet kerosene content is a fixed quantity for a given feed, this is formally the same optimisation problem as maximising the net production ($w_{C_8-C_{16}}^{\text{out}} - w_{C_8-C_{16}}^{\text{in}}$). The outlet fraction is reported as the primary indicator, with the inlet kerosene content explicitly recalled to avoid overstating the contribution of the hydrocracking step.

The second feed is the C_{17+} fraction of the same stream, obtained by removing all species lighter than C_{17} upstream of the reactor, whose composition is reported in Appendix A.4 leading to a weighted average molecular weight of 326.016 g/mol. This configuration reflects an alternative process layout in which only the heavy fraction, the fraction that genuinely requires conversion, is sent to the hydrocracker, while the lighter species are routed directly to the kerosene pool or to other upgrading units. Because this feed contains no C_8 - C_{16} at the inlet, maximising the outlet C_8 - C_{16} mass fraction is here equivalent to maximising the net kerosene production in the strict sense. Processing only the C_{17+} fraction also avoids exposing the kerosene already present in the full stream to overcracking, and confines the catalyst duty to the heavy paraffins, which is favourable both for product selectivity and for catalyst lifetime. It should nonetheless be noted that treating the inlet as a pure C_{17+} fraction is an idealising assumption: in practice, even with a dedicated upstream separation, some lighter species would inevitably remain in the stream entering the hydrocracker.

For both feeds, the conversion of the C_{17+} fraction, $X_{C_{17+}}$, is recorded alongside the C_8 - C_{16} yield as a secondary performance indicator, since it is the quantity most directly comparable with the conversion levels reported in the hydrocracking literature.

A central limitation of this work

Before describing the search procedure, one limitation must be emphasised, as it conditions the interpretation of every result that follows. The kinetic parameters of the model were estimated against an experimental dataset obtained with a specific C_4 - C_{70} n-paraffin distribution. The ULiège inlet stream characterised in Section 2.1 differs from this distribution, both in carbon-number coverage and in concentration profile. Although the underlying kinetic mechanism (hydroisomerization followed by cracking, with carbon-number-dependent rate constants) is expected to remain valid for any n-paraffin feed, the calibrated parameters were not adjusted to the present feedstock due to a lack of experimental data for the ULiège flow. Consequently, the operating conditions identified by the search must be regarded as indicative rather than absolute, and the absolute kerosene yield reported in Chapter 3 is more representative of an order-of-magnitude estimate than of a precise design value. The procedure described below should therefore be understood as a methodology that can be re-applied once the kinetic parameters have been refitted to the actual feedstock using experimental measures.

Thermodynamic validity condition

A further constraint is imposed by the thermodynamic condition established in Section 2.3.3: the overall vapour mole fraction of the reactive mixture must remain above 96% for the partial-pressure assumption underlying the kinetic model to hold. Because this condition cannot be evaluated within the Python simulation environment, it is verified post-hoc in Aspen Plus on the operating conditions identified by the search.

2.4.2 Search strategy for the kerosene-maximising operating conditions

The objective of this section is to identify, among all operating conditions (T , P , WHSV, H_2/wax) contained in the chosen domain, the combination that yields the highest $C_8\text{-}C_{16}$ mass fraction at the reactor outlet. Strictly speaking, this is not an optimisation problem: no closed-form objective function relates the four operating parameters to the kerosene yield, and the only way to evaluate the yield at a given operating point is to run the full reactor simulation. The procedure adopted here is therefore better described as a structured search.

A purely grid-based search would in principle suffice, but it is not optimal for one reason. Preliminary simulations show that the $C_8\text{-}C_{16}$ yield does not exhibit a single sharp peak; instead, several operating points produce essentially the same kerosene yield, within the numerical precision of the simulation. Selecting a single grid point as “the” best operating point would therefore be misleading, since a small refinement of the grid resolution could replace it with a different, equally good point. The procedure must instead identify the region of nearly equivalent solutions and apply an additional criterion to discriminate between them.

The search is organised in four phases, illustrated in Figure 2.5, and governed by three tolerances, Δ_{iso} , Δ_{best} and Δ_{dup} , whose role is defined phase by phase below, and whose numerical values define the two precision scenarios discussed at the end of the section.

Phase 1: Grid exploration

A grid of $9 \times 9 \times 9 \times 5 = 3645$ operating points is generated over the chosen domain: temperature, pressure and WHSV are each sampled at nine levels, and the H_2/wax ratio at five. This sampling is a compromise between computational cost and resolution: the H_2/wax axis is sampled most coarsely because the partial-pressure assumption leaves the model only weakly sensitive to this ratio, as confirmed by the sensitivity analysis of Section 3.2.2.3.1. At each grid point, the full reactor simulation is run with the backward differentiation formula (BDF) integrator, and the resulting $C_8\text{-}C_{16}$ mass fraction at the outlet is recorded. The BDF scheme is preferred because the material-balance system can become locally stiff in the high-temperature, low-WHSV corner of the domain; with this integrator, all 3645 points converge over both exploration domains, confirming that the model behaves consistently across the full operating window. The output of this phase is a four-dimensional response surface sampled on the grid.

Phase 2: Identification of the iso-yield set and selection of candidates

The grid points whose $C_8\text{-}C_{16}$ yield lies within Δ_{iso} of the global grid maximum form the *iso-yield set*. The tolerance Δ_{iso} is the first of the two parameters that define a precision scenario: a small Δ_{iso} retains only the points practically indistinguishable from the best one, whereas a larger value deliberately admits a broader region of nearly equivalent, or modestly suboptimal, operating conditions.

This set may contain anywhere from a few points to several hundred, scattered across the domain. Refining each of them with the local search of Phase 3 would be prohibitively expensive, so a small, representative subset of $N \leq 5$ points is extracted by *farthest-point sampling*. Farthest-point sampling consists of fixing a first point and then repeatedly adding the iso-yield point lying farthest from those already chosen. The four operating variables are normalised beforehand, so that each contributes equally to this distance regardless of its physical units and range. This spreads the N points across the iso-yield

set instead of letting them cluster in one corner. The first point is not chosen for distance but fixed as the iso-yield point with the lowest H_2/wax ratio. Because the final ranking (Phase 4) favours the lowest hydrogen consumption, seeding the subset this way guarantees that the most economical operating point is always carried forward into the refinement.

The spread matters because Phase 3 refines each point by a local search: from its starting point, the search climbs to a local optimum and stops there. Had the starting points been clustered together, they would all climb to the same optimum, and any other high-yield region located elsewhere in the domain would go unexplored. By spreading the starting points across the iso-yield set, each distinct high-yield region is likely to receive its own starting point and be explored in Phase 3.

Phase 3: Local refinement

Each of the N candidates is refined independently with the Nelder-Mead simplex method [89], clipped to the physical bounds of Table 2.3. Starting from the candidate, the search evaluates nearby operating points and moves towards higher $C_8\text{-}C_{16}$ yield, taking wider steps while it keeps improving and shorter ones as it approaches a maximum. It stops when no further gain is found. Each candidate yields one refined point, so Phase 3 produces at most N refined optima. This is also where the precision of the final result is set: the Phase 1 grid only locates promising regions, whereas the refinement varies all four parameters continuously and is therefore not limited by the grid spacing.

Refinement makes each candidate at least as good as its own starting point, but the N results need not reach the same yield: each climbs only as high as the region around its starting point allows, and because the candidates were deliberately spread out, some regions are higher than others. Two relationships between the refined points are therefore possible, and distinguishing them is left to Phase 4. Either two candidates reach genuinely different optima, at different yields; or they converge to the same optimum and differ only because each simplex halts within a finite numerical tolerance of it, reporting slightly different coordinates for what is physically a single point.

Phase 4: Filtering, ranking and verification

The N refined points are now compared with one another. The highest yield among them is taken as the reference, and any point lying more than Δ_{best} below this reference is discarded, having converged to a poorer part of the domain. This is what makes Δ_{best} useful even though the search started from N candidates: the N points generally end at different yields, and Δ_{best} keeps only those that reach the highest one, to within the tolerance. Δ_{best} can be set tighter than Δ_{iso} because the simplex locates the optimum more precisely than the coarse grid. The selected points are then deduplicated: two candidates that started far apart may converge to the same optimum; because each simplex stops within a finite numerical tolerance of it, the two are reported at slightly different coordinates rather than as a single point. Any two points closer than Δ_{dup} in the normalised parameter space are therefore treated as identical, merged, and reported once. Through these two steps the count falls from at most N to the set of genuinely distinct best points.

The remaining optima are ranked according to an economic criterion, which the user is free to define. In this work, hydrogen consumption is selected as the basis for the ranking, because hydrogen production by water electrolysis is the most energy-intensive step of the upstream Power-to-Kerosene chain (Section 1.1.2.2) and therefore assumed to dominate the operating cost of the integrated process. Among the four operating parameters, H_2/wax is the one whose minimisation translates most directly into operating-cost savings; among the others, pressure has only a marginal effect on operating cost since compressing the liquid wax feed requires negligible work, while temperature and WHSV act primarily on reactor sizing and heat integration. The criterion can be replaced or complemented by any other scalar function of the operating parameters, for instance, a weighted norm or a cost function incorporating capital-expenditure terms, without altering the rest of the procedure. If only one refined optimum survives the filtering step, the criterion plays no role.

The most economically attractive refined optimum is designated the recommended optimum. Two verifications are then carried out. First, a set of two-dimensional sensitivity slices of the C_8 - C_{16} yield is computed around the recommended optimum, by sweeping each pair of operating parameters while holding the other two fixed at their optimal values; these slices visualise the local response surface and confirm that the identified point lies in a smooth, robust region rather than on a numerical artefact. Second, the thermodynamic validity of the point is checked: the overall vapour mole fraction (VFRAC) of the inlet mixture is computed in Aspen Plus with the SRK equation of state, as described in Section 2.3.3. If it falls below 96%, the partial-pressure assumption underlying the kinetic model is no longer admissible, and the next-ranked optimum is examined instead.

The recommended operating conditions, the corresponding outlet composition and the verifications mentioned in this section are reported and discussed in Chapter 3.

Two precision scenarios

The procedure above is run under two distinct settings of the tolerances, answering two different questions.

In the strict-optimality scenario ($\Delta_{\text{iso}} = 0.05$ wt-%, $\Delta_{\text{best}} = 0.01$ wt-%), only operating points whose yield is indistinguishable from the global maximum, within the numerical precision of the simulation, are retained. The procedure then returns the single best operating point, or the very small set of genuinely equivalent points among which the economic criterion acts merely as a tie-breaker. This scenario answers the question: *what is the highest achievable C_8 - C_{16} yield, and at which conditions?*

In the relaxed-optimality scenario ($\Delta_{\text{iso}} = 5$ wt-%, $\Delta_{\text{best}} = 3$ wt-%), operating points whose yield lies up to a few weight-percent below the best are deliberately admitted. This widens the iso-yield set by orders of magnitude and transforms the economic criterion from a tie-breaker into a genuine selection lever. This scenario answers a different question: *how far can the operating conditions be relaxed, in particular towards a lower H_2 /wax ratio, if a small, deliberately accepted penalty on the kerosene yield is tolerated?*

The deduplication tolerance is kept at $\Delta_{\text{dup}} = 0.02$ (2% of the normalised parameter range) in both scenarios, as it concerns only the merging of coincident simplex endpoints and is independent of how strictly optimality is defined.

Combining the two feed streams, the two exploration domains and the two precision scenarios yields the results of eight searches presented in Chapter 3.




Figure 2.5: Four-phase search procedure used to identify the recommended operating point.

2.5 Reactor sizing model for LTFT wax hydrocracking

This section establishes the methodology used to derive the optimal operating conditions identified in Section 2.4.2 into the physical dimensions of the hydrocracking reactor.

The bench-scale unit on which the experiments of Pellegrini et al. were carried out is a trickle bed. In their kinetic study, however, this reactor is idealised as a plug-flow reactor (PFR) operating under isothermal conditions: axial and radial dispersion are neglected, the system is assumed perfectly homogeneous over any cross-section, and no mass-transfer limitation is retained, neither at the gas-liquid interface nor within the catalyst particles. This idealisation is justified, at laboratory scale, by the use of a finely powdered catalyst. Real trickle beds, by contrast, exhibit far more complex hydrodynamics (partial wetting, liquid maldistribution, axial and radial dispersion, wall channelling) and the PFR description holds only when the bed geometry sufficiently attenuates these effects, which is precisely the purpose of the geometric criteria imposed in Section 2.5.3. The sizing presented below should therefore be regarded as indicative rather than definitive: a construction-grade design would require several additional calculations that fall outside the scope of this work.

The sizing follows a sequential methodology. The optimal operating conditions fix the weight hourly space velocity which, together with the inlet wax mass flow rate, determines the required catalyst mass m_{cat} (Section 2.5.1). The catalyst mass and the bulk density of the bed then fix the packed-bed volume through $V_R = m_{\text{cat}}/\rho_{\text{bulk}}$ (Section 2.5.2), from which the reactor geometry, internal diameter D and bed height H , is finally derived under prescribed hydrodynamic constraints (Section 2.5.3).

A defining feature of this methodology is that the catalytic bed volume is governed exclusively by the catalyst mass. As recalled in Section 1.4.5.2, the random packing of solid particles leaves an inter-particle void fraction ε_b that is an intrinsic geometric feature of the bed, independent of any fluid flowing through it, and already embedded in the bulk density of the bed ρ_{bulk} through the relation $\rho_{\text{bulk}} = (1 - \varepsilon_b)\rho_p$, where ρ_p is the density of the catalyst particles. In operation, the liquid wax wets and trickles over the external surface of the pellets, while the hydrogen gas percolates through these pre-existing voids. Hydrogen therefore requires no volume to be specifically allocated to it: although it plays a central kinetic role through its partial pressure in the rate expressions, the hydrogen flow rate is not a sizing variable for the reactor. It is, however, a major determinant of the pressure drop across the bed, and for this reason re-enters the geometry discussion of Section 2.5.3.

2.5.1 Catalyst mass from WHSV

The weight hourly space velocity (WHSV) was adopted as the main sizing variable, as it directly relates the inlet wax mass flow rate to the catalyst mass. It is defined as

$$\text{WHSV} = \frac{\dot{m}_{\text{wax},\text{in}}}{m_{\text{cat}}} \quad (2.39)$$

where $\dot{m}_{\text{wax},\text{in}}$ is the inlet mass flow rate of LTFT wax and m_{cat} is the catalyst mass. Rearranging Eq. (2.39) yields the catalyst mass required to operate at a prescribed WHSV:

$$m_{\text{cat}} = \frac{\dot{m}_{\text{wax},\text{in}}}{\text{WHSV}} \quad (2.40)$$

In this formulation, WHSV acts as an inverse measure of contact time τ : lower WHSV values imply larger catalyst masses and longer contact times within the catalytic bed.

2.5.2 Packed-bed volume

The required packed-bed volume is obtained from the catalyst mass and the bulk density of the catalyst bed. The latter, whose value is difficult to find in the literature, was estimated from the reference packed bed of known catalyst mass and bed geometry used by Pellegrini et al. [4]. Assuming a cylindrical bed, the reference bed volume is expressed as

$$V_{\text{bed,ref}} = \frac{\pi D_{\text{ref}}^2 H_{\text{ref}}}{4} \quad (2.41)$$

where D_{ref} and H_{ref} are respectively the internal diameter and the bed height of the reference reactor. The corresponding bulk density is then

$$\rho_{\text{bulk}} = \frac{m_{\text{cat,ref}}}{V_{\text{bed,ref}}} \quad (2.42)$$

where $m_{\text{cat,ref}}$ is the catalyst mass in the reference bed.

Since the total bed volume comprises both the volume of the catalyst pellets and the inter-particle void space, the bulk density is lower than the particle density ρ_p ; the two quantities are related through the bed void fraction ε_b by

$$\rho_{\text{bulk}} = (1 - \varepsilon_b) \rho_p \quad (2.43)$$

Once ρ_{bulk} is known, the packed-bed volume required for the design catalyst mass is calculated as

$$V_R = \frac{m_{\text{cat}}}{\rho_{\text{bulk}}}. \quad (2.44)$$

2.5.3 Reactor geometry

The reactor is assumed to be a cylindrical packed bed, so that its volume relates the internal diameter D and the bed height H through

$$V_R = \frac{\pi D^2 H}{4} \quad (2.45)$$

Because the bed volume V_R is already fixed by the catalyst mass (Section 2.5.2), Equation (2.45) leaves the geometry with a single degree of freedom: selecting the diameter immediately fixes the height through $H = 4V_R/(\pi D^2)$. Diameter and height cannot therefore be chosen independently: at constant V_R , increasing the diameter necessarily shortens the bed, and reducing it lengthens the bed.

Two geometric criteria constrain this single degree of freedom, both ensuring that the packed bed behaves as an ideal plug-flow reactor. The first is a lower bound on the tube-to-particle diameter ratio, $D/d_p \geq 25$, which limits wall channelling and fluid bypassing along the reactor wall [91]. The second is a lower bound on the aspect ratio, $H/D \geq 5$, a standard design guideline that ensures the bed is sufficiently elongated for axial dispersion to remain negligible [92].

The geometry is determined by a progressive procedure starting from the catalyst particle size d_p . The wall-effect criterion sets the smallest admissible diameter, $D_{\text{min}} = 25 d_p$. At this diameter the bed is at its most slender: since V_R is fixed, the height is maximal and the aspect ratio $H/D = 4V_R/(\pi D^3)$ reaches its highest value. If this value exceeds 5, the reactor is excessively elongated. Such a slender geometry is penalised in terms of pressure drop: a small diameter produces simultaneously a longer bed and a higher superficial velocity of the fluids, the volumetric flow rate, and in particular that of the hydrogen-rich gas phase, being divided by a smaller cross-sectional area, and both effects increase the pressure drop across the packing, as described by the Ergun equation [93]. The diameter is therefore progressively increased above D_{min} ; at constant V_R this shortens the bed and lowers the aspect ratio, until $H/D = 5$ is reached. This defines the design geometry.

The value $H/D = 5$ is thus retained as the minimum aspect ratio compatible with plug-flow behaviour: it is the shortest bed for which axial dispersion and wall channelling remain negligible. Among all geometries satisfying both criteria, it is the one with the largest diameter and the shortest bed, and therefore the one of lowest pressure drop. A smaller ratio would reduce the pressure drop further, but would no longer guarantee plug-flow behaviour; a larger one would only increase the pressure drop without any compensating benefit. Because the retained diameter exceeds D_{min} , the wall-effect criterion $D/d_p \geq 25$ is satisfied a posteriori with a comfortable margin.

The pressure drop itself is not quantified in the present work; a full evaluation through the Ergun equation is left to detailed design. One simplifying assumption of the present sizing methodology should be made explicit. The catalyst particle size is taken to be that of the powdered catalyst used in the experimental campaign of Pellegrini et al. ($d_p=0.625$ mm), which was specifically chosen to suppress diffusional resistances during kinetic parameter estimation. In an industrial or pilot trickle bed, such fine particles would generate prohibitive pressure drops, since the resistance to fluid flow through a packed bed grows rapidly when particle size decreases; the catalyst would instead be shaped into pellets or extrudates of larger characteristic size, at the price of reintroducing intraparticle and interphase mass-transfer resistances. Both refinements lie beyond the scope of the present sizing methodology.

This procedure provides a consistent basis for deriving a reactor geometry that satisfies the hydrodynamic criteria required for plug-flow behaviour while keeping the catalytic bed compact.

2.5.4 Hydrogen co-feed

The hydrogen feed rate is fixed by the hydrogen-to-wax mass ratio retained among the optimal operating conditions, according to

$$\dot{m}_{H_2} = \left(\frac{H_2}{wax} \right) \dot{m}_{wax,in} \quad (2.46)$$

As established in the introduction of this section, the hydrogen flow rate is not a sizing variable for the reactor, since hydrogen flows through the inter-particle voids already accounted for in the bed volume. It nonetheless remains an essential operating specification: it sustains the partial pressure that governs the LHHW rate expressions, maintains the hydrocracking environment, and conditions the gas–liquid hydrodynamics of the trickle bed. In the present sizing procedure, the hydrogen-to-wax ratio is therefore treated as an operating input rather than as a geometric sizing variable.

2.5.5 Heat management

The thermal behaviour of the hydrocracking reactor results from the partial mutual compensation of two opposite contributions. The cracking reactions are endothermic, whereas the accompanying hydrogenation reactions are exothermic, in particular the saturation of the olefins contained in the Co-LTFT feed. Although olefin hydrogenation has been treated as a separate stoichiometric step for modelling purposes (Section 2.1.3), in the reactor as it would actually be operated these olefins are saturated in situ, within the hydrocracking bed itself. The heat released by this exothermic saturation offsets a substantial part of the endothermic cracking duty, so that hydrocracking of Fischer–Tropsch wax proceeds under near-isothermal conditions [11].

As a direct consequence, and in agreement with the work of A. de Klerk (2011), no quench streams or interbed cooling are required for FT wax hydrocracking. The engineering design must, however, make sufficient provision for preheating the feed, so that the reactor can be brought to and operated across the temperature range of the catalyst [11]. The near-isothermal character of the process also provides an a posteriori justification for the isothermal plug-flow reactor model adopted in this work: heat management reduces, in practice, to adequate feed preheating rather than to in-bed temperature control.

Chapter 3

Results and discussions

Chapter 2 constructed the hydrocracking reactor model in three parts: the component-level kinetic model built on the LHHW formalism of Pellegrini et al. (2007) [3], the structured search procedure devised to locate the kerosene-maximising operating conditions, and the sizing methodology deriving those conditions into a reactor geometry. That chapter was deliberately confined to the construction of the model; nothing was said of whether it reproduces experimental reality, of what it predicts for the ULiège feed, or of how those predictions should be interpreted. The present chapter addresses these three questions in turn.

Their order is not arbitrary. A predictive model carries no weight until it has been confronted with experimental data, and the operating conditions it recommends become meaningful only once the model that produced them has been validated. The chapter therefore proceeds from validation to application, and from application to a broader discussion of the operating space, in three sections. Section 3.1 establishes the validation: since no bench-scale hydrocracking unit is available at ULiège, the model is confronted with the experimental dataset of Pellegrini et al. (2007) [3], which requires recovering the inlet distribution that the original publication omits, demonstrating that the simplified stoichiometric and thermodynamic-equilibrium descriptions are structurally unable to reproduce the experimental behaviour, and calibrating the LHHW kinetic parameters, themselves absent from the publication, against the published results. Section 3.2 applies the validated model to the feedstock for which the unit is intended, the Fischer-Tropsch hydrocarbon stream of the ULiège Power-to-Kerosene process: the operating point maximising the C₈-C₁₆ yield is identified, the kinetic profiles producing it are analysed, the robustness of that point is probed through a sensitivity study, and the corresponding reactor geometry is derived. Finally, Section 3.3 broadens the analysis beyond a single recommended point by applying the search procedure to the eight cases defined in Chapter 2: two feed streams, two exploration domains and two precision scenarios, so that the influence of feed selection, of extrapolation beyond the calibration range, and of the economic ranking criterion can each be assessed independently.

It should be emphasised, as a caveat conditioning the reading of the entire chapter, that the kinetic parameters are calibrated on the C₄-C₇₀ feedstock of Pellegrini et al. and are not refitted to the ULiège stream, for which no experimental data are available. The yields and operating conditions reported below should therefore be understood as order-of-magnitude indications, and as a methodology to be re-applied once feed-specific data exist, rather than as definitive design values. The extent to which this calibration limitation affects the conclusions is assessed in Section 3.3.4, where the search is repeated on the Pellegrini calibration feed. The limitations inherited from the underlying model, primarily the midpoint cracking stoichiometry and the partial-pressure assumption, are likewise made explicit at each point where they bear on the interpretation of the results.

3.1 Kinetic model validation

The validation of any reactor model requires experimental data against which its predictions can be assessed. Since the University of Liège does not operate a bench-scale LTFT wax hydrocracking unit, no experimental dataset was available internally for model validation. Therefore, as discussed in the previous chapter, the Pellegrini et al. (2007) [3] article presenting the all-components model, in which the model predictions are compared with experimental results, was selected as the reference model. However, exploiting this reference requires a preliminary step: the article does not report the inlet composition used to generate its published results, so that no comparison between an implemented model and the published curves is meaningful until the inlet is identified. The identification of this inlet distribution is therefore addressed first in Section 3.1.1, before any model comparison is undertaken.


3.1.1 Identification of the Pellegrini (2007) inlet distribution

As explained in Section 2.2.2, the article of Pellegrini et al. (2007) is structured in two successive parts: in the first, the kinetic constants take a single fixed value per lump; in the second, these lump values are fitted to continuous functions of carbon number, and a non-linear regression is performed to obtain the final optimised coefficients. In neither part is the inlet distribution reported, only the model-predicted outlet distributions and the experimental data points are given. This omission represents a practical obstacle to model reproduction: without knowledge of the inlet, it is not possible to verify that an implementation of the model produces results consistent with those published. A request to the corresponding author for the original data received no response, so the inlet had to be recovered by hypothesis testing.

The strategy adopted relies on the first (lumped-constant) part, whose kinetic constants are fully reported (Appendix A.3), unlike those of the second part; the inlet distribution is therefore the only unknown of that model. Each candidate inlet is used as input to the lumped model, and the resulting outlet is compared with the published lumped-model curves rather than with the experimental data points. The model curves are used as the reference because the inlet is recovered through Pellegrini's model itself: comparing against the experimental points would require the inlet to compensate for the model's own deviation from experiment, which would bias the identification. Agreement with the model curves therefore provides a direct measure of the validity of each candidate inlet.

First hypothesis: lumped inlet from Pellegrini et al. (2004)

The natural first candidate is the feedstock described in the preceding article of the same series, Pellegrini et al. (2004) [5], which reports the inlet composition as four carbon-number cuts with their associated weight fractions, reproduced in Figure 3.1. Since this lumped description does not resolve the individual species within each cut, a uniform molar distribution is assumed within each lump.



Cut (carbon atoms)	Weight fraction
C ₉ –	0.037
C ₁₀ –14	0.176
C ₁₅ –22	0.301
C ₂₂ +	0.486

Figure 3.1: Inlet distribution constructed from the lumped feedstock composition reported by Pellegrini et al. (2004) [5], assuming a uniform mass distribution of species within each carbon-number cut.

The outlet distribution simulated with this inlet is shown in Figure 3.2. The comparison with the lumped-model curves published by Pellegrini et al. (2007) reveals a systematic mismatch, quantified in the parity plot of Figure 3.3: sharp discontinuities appear at the lump boundaries, producing step-like transitions in the carbon-number profiles that are entirely absent from the published distributions. These discontinuities arise directly from the uniform-distribution assumption within each lump. All species sharing the same lump necessarily carry the same initial mass fraction, so that the molar inlet profile presents abrupt jumps at the lump boundaries, which propagate into the outlet. Since the published curves are smooth, this behaviour is inconsistent with the inlet used by Pellegrini et al. (2004), and the first hypothesis is rejected.




Figure 3.2: Outlet carbon-number distributions simulated with the lumped 2004 inlet compared with the lumped-model curves published by Pellegrini et al. (2007) [3], under conditions ($T = 367^\circ\text{C}$, $P = 53.75$ bar, $H_2/\text{wax} = 0.1275$ kg/kg, $\text{WHSV} = 2.5$ kg_{n-C},in kg_{cat}⁻¹ h⁻¹).




Figure 3.3: Parity plots of the outlet molar fractions simulated with the component-level 2004 inlet against the lumped-model curves of Pellegrini et al. (2007) [3,5].

Second hypothesis: component-level inlet from Pellegrini et al. (2008)

The smooth character of the published outlet distributions implies that the inlet was already specified at the component level, with individual mass fractions assigned to each n-paraffin rather than uniform values within broad cuts. The next article in the series, Pellegrini et al. (2008) [4], reports precisely such a distribution as a graph in the article: a component-resolved C_4 - C_{70} n-paraffin inlet, reported after digitization in Appendix A.5 and illustrated in Figure 3.4. The hypothesis that this distribution was used in the 2007 article is supported by the fact that Pellegrini et al. (2008) explicitly state that their model extends the 2007 formulation to account for vapour-liquid equilibrium, implying that the same inlet composition serves as the calibration basis throughout the series. However, since the final distributions reported in the two articles were not obtained under the same operating conditions, a direct comparison of the experimental results did not allow this conclusion to be drawn straightforwardly.




Figure 3.4: n-Paraffin inlet distribution reported by Pellegrini et al. (2008) [4].

The outlet distribution obtained with this inlet is reported in Figure 3.5. The agreement with the lumped-model curves of Pellegrini et al. (2007) is markedly improved: the step-like artefacts disappear entirely, and the simulated profiles follow the same smooth variation with carbon number as the published model. The agreement is quantified in the parity plots of Figure 3.6, which report $R^2 = 0.9981$ for n-paraffins and $R^2 = 0.9992$ for iso-paraffins. These values approach the upper limit of what this comparison method can resolve: the curves from the articles were obtained by digitisation of the published figures, and any residual scatter reflects digitization noise rather than a genuine model discrepancy. The Pellegrini et al. (2008) inlet is therefore retained as the inlet distribution for all subsequent comparisons involving the Pellegrini et al. (2007) model.




Figure 3.5: Outlet carbon-number distributions simulated with the component-level 2008 inlet compared with the lumped-model curves published by Pellegrini et al. (2007) [3], under conditions ($T = 367^\circ\text{C}$, $P = 53.75\text{ bar}$, $\text{H}_2/\text{wax} = 0.1275\text{ kg/kg}$, $\text{WHSV} = 2.5\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$).

A clarification should be made regarding the graphical representation shown in Figure 3.5. Following the convention adopted by Pellegrini et al. (2004), the C_4 isomer was grouped with the corresponding n-paraffin for graphical representation purposes in the all-components model as well [3, 5]. This grouping does not affect the calculations, since iso- C_4 and n- C_4 were treated separately in the material balance equations detailed in Section 2.3.1.




Figure 3.6: Parity plots of the outlet molar fractions simulated with the component-level 2008 inlet against the lumped-model curves of Pellegrini et al. (2007) [3, 4].


3.1.2 Stoichiometric model

In the literature, when a hydrocracking unit is simulated as part of a broader process chain, it is common practice to represent the reactor as a stoichiometric unit, on the grounds that process-level performance metrics can be evaluated without resolving the detailed reaction kinetics. Vaquerizo and Rego-Fernández [94], for instance, adopt this approach in their power-to-kerosene process model, noting that a stoichiometric reactor allows straightforward integration into a process simulator while remaining ac-


ceptable when the objective is the evaluation of overall process performance rather than the detailed reactor behaviour. This section examines whether this level of description is sufficient for the present work, whose objective is to reproduce the experimental data as closely as possible and to identify the operating conditions that maximise the C_8 - C_{16} yield.

Two stoichiometric cracking models are considered, corresponding to the two product-distribution assumptions introduced in Section 1.4.2: the midpoint approximation and the breakage probability function. In both cases, all species, heavier than C_4 for the midpoint approximation and heavier than C_8 for the breakage probability cracking, are assumed to react completely. The special cracking stoichiometry applied to iso- C_5 , Reaction (2.9), is retained for the midpoint approximation stoichiometry model following the convention adopted by Pellegrini et al. (2007) [3]. For each variant, two scenarios are simulated: primary cracking only (one cracking pass), and primary cracking followed by secondary cracking (two successive passes). Secondary cracking is modelled by applying the primary-cracking step a second time, the outlet of the first pass being used as the feed to the second. It should be noted that the stoichiometric model does not incorporate any reaction mechanism and therefore makes no distinction between n-paraffins and iso-paraffins: the outlet is described solely as a redistribution of total paraffinic moles across carbon numbers. The comparison with the experimental data of Pellegrini et al. (2007) [3] is consequently carried out on a total carbon-number basis, independently of the molecule's branching.

The results are shown in Figures 3.7a to 3.8b, where the simulated outlet distributions are compared with the experimental outlet distribution of Pellegrini et al. (2007) under conditions ($T = 367^\circ\text{C}$, $P = 53.75$ bar, $H_2/\text{wax} = 0.1275$ kg/kg, $\text{WHSV} = 2.5$ kg_{n-C₁₀}/in kg_{cat}⁻¹ h⁻¹). Neither the midpoint approximation nor the breakage probability function, whether applied over one or two cracking passes, reproduces the shape of the experimental distribution. The midpoint model concentrates cracking products in a narrow region around half the parent carbon number, producing a distribution that bears little resemblance to the broad, smoothly decaying experimental profile. The breakage probability function distributes products more evenly across carbon numbers, which improves the qualitative shape, yet the simulated distribution remains incompatible with the experimental data. Applying a second cracking pass shifts additional mass toward lighter species, moving the simulated distribution further from, rather than closer to, the experimental reference.



(a) Midpoint approximation.



(b) Breakage probability function.

Figure 3.7: Comparison between the simulated outlet carbon-number distribution and the experimental data of Pellegrini et al. (2007) [3] after one stoichiometric cracking pass with complete conversion of all C_{5+} species.




Figure 3.8: Comparison between the simulated outlet carbon-number distribution and the experimental data of Pellegrini et al. (2007) [3] after two successive stoichiometric cracking passes with complete conversion at each pass.

Beyond the inability to reproduce the experimental distribution, the stoichiometric model suffers from a more fundamental limitation with respect to the objectives of this work. By construction, the outlet distribution of a stoichiometric reactor depends only on the inlet composition and the assumed cracking stoichiometry. This approach remains limited, as it does not account for reaction kinetics, catalyst effects, or the influence of operating conditions. Two simulations run at different operating conditions but with the same inlet and cracking stoichiometry produce identical outlet distributions. This property is incompatible with the objective of identifying the operating conditions that maximise the kerosene yield.

3.1.3 Thermodynamic model

Following the same logic as for the stoichiometric model, a thermodynamic approach was investigated as a second simplified alternative, to determine whether such a description is sufficient or whether a full kinetic model is required. Two reactor models were built in Aspen Plus: a Gibbs reactor (RGibbs), which minimises the total Gibbs free energy of the system without any prescribed reaction network, and an equilibrium reactor (REquil), which enforces thermodynamic equilibrium for a specified set of reactions. The REquil model was run with both the midpoint and the breakage probability cracking stoichiometries introduced in Section 1.4.2. All simulations were carried out under the same operating conditions as the stoichiometric model ($T = 367^\circ\text{C}$, $P = 53.75$ bar, $\text{H}_2/\text{wax} = 0.1275$ kg/kg), noting that, like the stoichiometric reactor, a thermodynamic equilibrium model carries no dependence on the weight hourly space velocity, since the outlet state is determined by thermodynamics alone rather than by contact time. In addition, no distinction between *n*-paraffins and iso-paraffins is made either.

Before presenting the results, an important distinction must be drawn with respect to the model of Pellegrini et al. (2008) [4], which incorporates vapour-liquid equilibrium calculations and achieves improved agreement with experimental data compared to its predecessor. This improvement does not arise from treating the reactor as a thermodynamic equilibrium unit. The 2008 model retains the same isothermal plug-flow reactor and the same LHHW kinetic rate expressions as the 2007 model; vapour-liquid equilibrium enters solely through the replacement of partial pressures by fugacities in the rate expressions, so that the phase distribution of each species is correctly accounted for when evaluating the reaction driving force. The reactor remains kinetically controlled throughout; thermodynamics governs the phase split, not the reaction rate. The equilibrium reactors investigated in the present section are therefore of a fundamentally different nature.

Due to the limitations of the Aspen Plus thermodynamic database for mass balances with pseudo-components beyond C_{50} , the RGibbs and REquil simulations were restricted to species up to C_{30+} , following the representation adopted by doctoral researchers at ULiège.

The results obtained with the RGibbs model are shown in Figure 3.9. The outlet distribution collapses entirely toward methane and ethane : the molar fraction of hydrogen drops from 0.952 at the inlet to 0.105 at the outlet, indicating near-complete consumption. At high temperatures and pressures, the Gibbs free energy of the system is minimised by decomposing long-chain paraffins into the lightest possible alkanes which represent the most stable configuration under these conditions. The resulting distribution bears no resemblance to the experimental outlet, confirming that the reactor cannot be modelled as an equilibrium reactor.




Figure 3.9: Outlet carbon-number distribution predicted by the RGibbs model (logarithmic scale) compared with the initial distribution and the experimental outlet of Pellegrini et al. (2007) [3], at $T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $\text{H}_2/\text{wax} = 0.1275\text{ kg/kg}$.

The REquil model yields qualitatively similar conclusions. In this case too, the equilibrium reactor predicts a distribution heavily shifted toward light species. With the midpoint cracking stoichiometry, the hydrogen molar fraction decreases from 0.952 to 0.745; with the breakage probability function, it decreases to 0.836. In both cases, the predicted outlet distributions shown in Figures 3.10 and 3.11 remain incompatible with the experimental reference.




Figure 3.10: Outlet carbon-number distribution predicted by the REquil model with midpoint cracking stoichiometry (logarithmic scale) compared with the initial distribution and the experimental outlet of Pellegrini et al. (2007) [3], at $T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $\text{H}_2/\text{wax} = 0.1275\text{ kg/kg}$.




Figure 3.11: Outlet carbon-number distribution predicted by the REquil model with the breakage probability cracking stoichiometry (logarithmic scale) compared with the initial distribution and the experimental outlet of Pellegrini et al. (2007) [3], at $T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $\text{H}_2/\text{wax} = 0.1275\text{ kg/kg}$.

Taken together, the results of both the stoichiometric and thermodynamic approaches demonstrate that neither is capable of reproducing the experimental outlet distribution or of predicting the effect of operating conditions on the product yield. The stoichiometric model is insensitive to all four operating parameters by construction; the thermodynamic model responds to temperature, pressure and hydrogen/wax ratio and predicts a degree of conversion that is physically unrealistic, confirming that the hydrocracking reaction under the conditions studied is kinetically controlled rather than thermodynamically limited. A kinetic model is therefore necessary. The LHHW formulation selected in Section 2.2.2, in which reaction rates depend explicitly on temperature, partial pressures, and chain length, provides the required sensitivity and constitutes the basis of all subsequent simulations.

3.1.4 LHHW kinetic model: parameter estimation and model reproduction

As established in Sections 3.1.2 and 3.1.3, since neither the stoichiometric nor the thermodynamic approach satisfactorily reproduces the experimental outlet distribution, a kinetic description is required. As the reference model for this work, the LHHW model of Pellegrini et al. (2007) [3], implemented at the component level as described in Section 2.3, constitutes the basis of all subsequent simulations. Before it can be applied to the ULiège inlet stream, the final expressions of the kinetic parameters used by the authors but not reported in the article, must be found by identifying the 15 coefficients appearing in the functional forms of Equations (2.28)–(2.34) following the dedicated parameter estimation procedure described in Section 2.3.2. Two estimation targets are considered successively: the model curves and the experimental data published by Pellegrini et al. (2007) in their Figures 2.3a and 2.3b. The results of the two approaches are presented and compared below.

Estimation against the Pellegrini model curves

The objective of this first estimation is to reproduce the outlet distributions modelled by Pellegrini et al. (2007) as faithfully as possible, so that all subsequent simulations are anchored to a known, validated reference. The regularised optimisation described in Section 2.3.2 was applied with the Pellegrini model curves as the target, digitised under conditions ($T = 367\text{ }^{\circ}\text{C}$, $P = 53.75\text{ bar}$, $\text{H}_2/\text{wax} = 0.1275\text{ kg/kg}$, $\text{WHSV} = 2.5\text{ kg}_{\text{n-C},\text{in}}/\text{kg}_{\text{cat}}^{-1}\text{ h}^{-1}$).

The regularised optimisation balances two competing demands: reproducing Pellegrini's reference curves as faithfully as possible, while keeping the kinetic parameters close to their published values, thereby

preserving their physical meaning. This second demand is enforced by the penalty term of Equation (2.36), which is zero when the parameters match the reference and grows as they drift away from it. At the start of the procedure, the parameters are exactly the reference expressions of Equations (2.21)-(2.27), so the penalty is inactive and the objective is set by the quality of the fit alone ($\mathcal{L}_{SSE} = 1.58 \times 10^{-3}$). Over the three successive stages, this fit improves markedly, the sum of squared errors dropping by about two orders of magnitude ($\mathcal{L}_{SSE} = 1.44 \times 10^{-5}$). Crucially, this gain is obtained without the parameters straying far from the reference: the penalty reached at the optimum stays very small (1.74×10^{-7} , well below the fit term itself), which means the parameters have moved only marginally. The optimised model therefore reproduces the reference distributions while remaining physically consistent with the published parameters. The resulting optimised kinetic expressions are:

$$K_{L,n}(n_c) = 0.4347 \exp(0.3898 n) \quad (3.1)$$

$$K_{L,iso}(n_c) = 0.1078 \exp(0.3855 n) \quad (3.2)$$

$$K_{eq}(n_c) = 1.1268 n_c^2 - 3.1914 n_c + 3589.14 \quad (3.3)$$

$$k_{isom,0}(n_c) = 5.265 \times 10^{16} n^{8.021} \quad (3.4)$$

$$k_{cr,0}(n_c) = 1.981 \times 10^{18} n^{5.851} \quad (3.5)$$

$$E_{isom}(n_c) = 22226 \ln(n) + 119599 \quad (3.6)$$

$$E_{cr}(n_c) = 29440 \ln(n) + 87185 \quad (3.7)$$

These coefficients are reported with their full fitted precision rather than rounded. Several of them enter the rate expressions through an exponential or power-law dependence on the carbon number, which spans the entire C_1 - C_{70} range, and this strong non-linearity amplifies small relative variations: altering even the fourth or fifth significant figure of the power-law exponents, of the exponential coefficients or of the activation energies measurably shifts the predicted distribution. The remaining coefficients are kept at the same precision so that re-evaluating the expressions reproduces the calibrated curves exactly, without adding a rounding error on top of the fit.

The comparison between the simulated distributions before and after optimisation is shown in Figure 3.12. Before optimisation, the initial expressions of Pellegrini et al. produce distributions that deviate noticeably from the target curves: the n-paraffin profile underestimates the peak region while the shape and amplitude of the iso-paraffin distribution are both poorly captured. After optimisation, both profiles are reproduced with excellent agreement across the full C_4 - C_{70} carbon-number range, confirming that the regularised procedure successfully recovers a parameter set consistent with the published model.





Figure 3.12: Outlet molar fraction distributions of n-paraffins (left) and iso-paraffins (right) simulated before and after parameter optimisation against the model curves of Pellegrini et al. (2007) [3], under conditions ($T = 367$ °C, $P = 53.75$ bar, $H_2/wax = 0.1275$ kg/kg, $WHSV = 2.5$ kg_{n-C},in kg_{cat}⁻¹ h⁻¹).


The corresponding evolutions of the optimised kinetic constants as functions of carbon number are

shown in Figure 3.13 and detailed numerically in Appendix A.6. The relative deviations of the optimised coefficients from the initial reference values of Pellegrini et al. remain generally modest. The largest adjustment concerns the pre-exponential factor of the Langmuir adsorption constant for n-paraffins, $K_{L,n,act}$, which increases from 0.10 to 0.435, a relative deviation of +335%. Although the absolute change appears large, the exponential functional form is preserved and the resulting $K_{L,n}$ curves remain approximately within the same order of magnitude as the reference, as visible in Figure 3.13b. All other parameters deviate by less than 8% in relative terms, with the equilibrium constants K_{eq} and the activation energy expressions E_{isom} and E_{cr} remaining virtually unchanged, as enforced by the higher regularisation weights assigned to those parameters.


The physical interpretation of the optimised parameter set is consistent with the underlying reaction mechanism and with the observations reported by Pellegrini et al. [4, 5]. Both the pre-exponential factors $k_{isom,0}$ and $k_{cr,0}$ and the Langmuir adsorption constants $K_{L,n}$ and $K_{L,iso}$ increase with carbon number, reflecting the greater reactivity of heavier paraffin molecules. The activation energies for isomerization and cracking are of similar magnitude and both increase with chain length, in agreement with [4, 71]. The Langmuir constants for n-paraffins are greater than those for iso-paraffins across the full carbon-number range, reflecting the stronger adsorption of linear species on the catalyst surface [4]. Finally, the equilibrium constants K_{eq} increase with n_c , as expected from the growing multiplicity of branched isomers with increasing chain length.




(a) Activation energies E_{isom} and E_{cr} .



(b) Langmuir constants $K_{L,n}$ and $K_{L,iso}$.



(c) Pre-exponential frequency factors $k_{isom,0}$ and $k_{cr,0}$.



(d) Equilibrium constants K_{eq} .

Figure 3.13: Carbon-number dependence of the kinetic parameters before optimisation (initial expressions of Pellegrini et al. (2007)) and after optimisation targeting the Pellegrini model curves.

Estimation against the experimental data

The same regularised optimisation framework was then applied with the digitised experimental data of Pellegrini et al. (2007) as the target, following the same three-stage procedure. This conceptually replicates the authors' original regression and yields a second parameter set, defined by Equations (3.8) to (3.14), which can be compared with the one regressed on the model curves.

$$K_{L,n}(n_c) = 8.163 \cdot \exp(0.3894 n_c) \quad (3.8)$$

$$K_{L,\text{iso}}(n_c) = 0.1286 \cdot \exp(0.3901 n_c) \quad (3.9)$$

$$K_{\text{eq}}(n_c) = 0.08282 n_c^2 - 4.418 n_c + 7972 \quad (3.10)$$

$$k_{\text{isom},0}(n_c) = 5.793 \times 10^{16} \cdot n_c^{8.285} \quad (3.11)$$

$$k_{\text{cr},0}(n_c) = 5.314 \times 10^{17} \cdot n_c^{6.477} \quad (3.12)$$

$$E_{\text{isom}}(n_c) = 2.183 \times 10^4 \ln(n_c) + 1.176 \times 10^5 \quad (3.13)$$

$$E_{\text{cr}}(n_c) = 2.772 \times 10^4 \ln(n_c) + 9.585 \times 10^4 \quad (3.14)$$

Despite directly targeting the experimental data, the optimised model fails to achieve a satisfactory fit for both n-paraffins and iso-paraffins. Yet, at first sight, as shown in Figure 3.14, the outcome appears visually closer to the experimental data than the model curves of Pellegrini et al. [3], shown in Figure 3.12, and might even seem to improve upon their fit.




Figure 3.14: Outlet molar fraction distributions of n-paraffins (left) and iso-paraffins (right) obtained after parameter optimisation against the experimental data of Pellegrini et al. (2007) [3], under conditions ($T = 367 \text{ }^\circ\text{C}$, $P = 53.75 \text{ bar}$, $\text{H}_2/\text{wax} = 0.1275 \text{ kg/kg}$, $\text{WHSV} = 2.5 \text{ kg}_{\text{n-C},\text{in}} \text{ kg}_{\text{cat}}^{-1} \text{ h}^{-1}$).

This apparent improvement is misleading: as shown in Figure 3.15, the visual gain is obtained only at the cost of the physical consistency of the parameters, the equilibrium constant K_{eq} in particular losing its physical significance. The optimised set is therefore not a genuine improvement; the residual remains high and the parameters are not physically consistent.




Figure 3.15: Carbon-number dependence of the kinetic parameters before optimisation (initial expressions of Pellegrini et al. (2007)) and after optimisation targeting the experimental data.

The origin of this limitation is structural, and lies in the cracking stoichiometry rather than in the kinetic parameters. As discussed in Section 1.4.2.2.1, the midpoint approximation places the products of each cracking event at a single carbon number, or at the two adjacent carbon numbers for an odd-numbered parent, corresponding to half of the parent chain length. This fixes the production routes of the model: a product of carbon number i can only be fed by the cracking of the single parent $2i$, whereas in practice, and in the breakage probability formulation, it receives fragments from every heavier parent $j > i$. The light and middle iso-paraffins (C_5 – C_{15}), which in reality accumulate fragments from across the heavy tail, are therefore deprived of these production routes, which is the direct cause of their under-production. This product distribution is determined by the stoichiometric coupling terms of the balance equations and is fixed once the midpoint approximation is adopted: the 15 optimisation coefficients only scale the reaction rates, that is, how much and how fast isomerisation and cracking proceed. They cannot alter which fragments a given parent yields, and therefore cannot redistribute the cracked mass along the carbon-number axis. The production routes that would feed the light and middle iso-paraffins from every heavier parent other than $2i$ are structurally absent from the balance equations, and cannot be reinstated by any choice of parameters.

Faced with this constraint, the optimiser compensates through kinetic adjustments that lack physical grounding: K_{eq} departs significantly from the reference quadratic profile, adopting a near-flat shape, and the Langmuir constant for n-paraffins is overestimated. These distortions, visible in Figure 3.15, signal that the optimiser has reached the limits of what a structurally constrained model can achieve, not that a genuinely better parameter set has been identified.

The contrast with the estimation against the Pellegrini model curves is decisive here, and can be quantified. Using the same procedure, this fit drove the $\mathcal{L}_{SSE}$ down to $\mathcal{L}_{SSE} \approx 1.44 \times 10^{-5}$ against the Pellegrini model curves (Figure 3.12), whereas against the experimental data it could not be driven below $\mathcal{L}_{SSE} = 1.34 \times 10^{-3}$, roughly two orders of magnitude higher. This gap is informative. When the target curve is itself reproducible by the model, the $\mathcal{L}_{SSE}$ collapses to near zero, which confirms that the optimiser and the regularisation behave correctly.

This interpretation is consistent with the progression of the models reviewed in Section 1.4.3.3: the introduction of vapour-liquid equilibrium, which does not modify the cracking stoichiometry, corrected the H_2 /wax and conversion behaviour but explicitly retained the midpoint approximation and did not resolve the iso-paraffin distribution, whereas its replacement by the breakage probability function did. It is also consistent with Pellegrini et al. (2007) themselves, who noted that their optimised continuous-function model yields lower iso-paraffin values than the experimental data and explicitly identified the midpoint cracking assumption as the hypothesis to be relaxed in subsequent work [3,6]. The present results confirm that finding: targeting the experimental data directly cannot resolve the limitation, and the discrepancy between the predicted and experimental iso-paraffin distributions must be attributed to the model structure rather than to an insufficiency of the parameter estimation procedure.

Selection of the parameter set for subsequent simulations

The comparison between the two estimation targets leads to a clear conclusion. The inability to match the experimental iso-paraffin distribution is a structural limitation of the midpoint cracking stoichiometry, not a deficiency of the kinetic parameter set. Attempting to correct it by targeting the experimental data produces a distorted parameter set without achieving a genuine improvement in model quality. The parameter set obtained by targeting the Pellegrini model curves, Equations (3.1)–(3.7), is therefore retained for all subsequent simulations. It reproduces the output distributions of the Pellegrini et al. model with a $\mathcal{L}_{SSE}$ of 1.44×10^{-5} , while remaining physically consistent with the basic model, as confirmed by the small regularization penalty. The residual underestimation of iso-paraffins relative to the experimental data is an acknowledged limitation inherited from the midpoint approximation of Pellegrini et al. (2007), and is discussed further in the perspectives of Chapter 4.

3.2 Application to the ULiège inlet stream

The preceding section established that a kinetic model, selected as the LHHW model of Pellegrini et al. (2007) [3] in this work, is needed to reproduce the experimental hydrocracking behaviour and to respond to the four operating variables. It further retained the parameter set obtained against the published model curves as the reference kinetic description. This validated model is now applied to the feedstock for which the hydrocracking unit is ultimately intended: the Fischer-Tropsch hydrocarbon stream of the ULiège Power-to-Kerosene process. It must be recalled that the kinetic parameters were calibrated on the C_4 – C_{70} distribution of Pellegrini et al. and not refitted to the ULiège feed, so that the present application is, strictly speaking, an extrapolation whose quantitative outcomes should be interpreted with caution. The section proceeds in three steps: Section 3.2.1 characterises the stream entering the reactor after the olefin-hydrogenation pre-treatment and contrasts it with the Pellegrini calibration feedstock; Section 3.2.2 identifies the operating point that maximises the C_8 – C_{16} yield, analyses the kinetic profiles that produce it, and assesses its robustness through a sensitivity study; and Section 3.2.3 translates the recommended operating conditions into the physical dimensions of the reactor.

3.2.1 Inlet composition after olefin hydrogenation

As described in Section 2.1.3, the 3.69 mol% of olefins present in the FT hydrocarbon product stream of ULiège are hydrogenated to their corresponding n-paraffins with full conversion. The resulting stream, composed exclusively of n-paraffins C_nH_{2n+2} spanning carbon numbers C_1 to C_{70+} , constitutes the inlet

to the hydrocracking reactor for all simulations presented in this chapter. Its detailed molar composition is reported in Appendix A.7.

Figure 3.16 illustrates the contribution of each species class to the total inlet. The blue bars represent the n-paraffins originally present in the FT product stream, while the red bars represent the olefins prior to hydrogenation; each olefin C_nH_{2n} is converted to the n-paraffin of the same carbon number, so that the total after hydrogenation corresponds exactly to the sum of both. Since the olefin content is modest relative to the paraffin content, the shape of the total distribution after hydrogenation closely follows that of the original n-paraffin fraction, with the olefin contribution introducing a small but visible upward shift concentrated in the C_3 - C_{13} range. The distribution peaks at C_{15} , and decays monotonically for heavier species, reaching negligible fractions beyond C_{40} .




Figure 3.16: Molar composition of the FT hydrocarbon product stream of ULiège before and after olefin hydrogenation.

As shown in Figure 3.17, the inlet stream obtained after olefin hydrogenation differs in two important aspects from the feedstock used by Pellegrini et al. to calibrate the kinetic model. First, the ULiège stream contains light species from C_1 , originating from the cobalt-based Fischer-Tropsch reactor and present in non-negligible amounts at C_1 - C_4 , whereas the Pellegrini et al. (2008) inlet begins at C_4 with no species lighter than butane. Second, the two distributions differ in their heavy-end character. Although both distributions peak within a similar carbon-number region, the Pellegrini distribution exhibits a higher molar fraction in the C_{20} - C_{50} range, leading to a weighted-average molar mass of 273.2 g/mol, as opposed to 236.1 g/mol for the ULiège distribution. The ULiège distribution, by contrast, decays more steeply beyond C_{20} and has a higher proportion of species already falling within or just above the kerosene range (C_8 - C_{16}). This lighter overall character arises from the specific operating conditions of the upstream Fischer-Tropsch reactor at ULiège, which was optimised to target C_8 - C_{16} hydrocarbons, and from the flash-separation at 150 °C that conditions the heavy fraction entering the hydrocracking unit.

These compositional differences have a direct consequence for the reliability of the kinetic model applied to the ULiège stream: as stated in Section 3.1.4, the kinetic parameters were calibrated on the heavier Pellegrini feedstock, and applying them to the lighter ULiège distribution constitutes an extrapolation whose predictions must be interpreted with appropriate caution.

Regarding the thermodynamic-validity assumption of the model, the fact that the ULiège feed is lighter than that of Pellegrini is favourable, as it allows the ULiège inlet stream to remain within the validity domain of the model over a range comparable to, or even slightly wider than, that obtained for the Pellegrini inlet. For example, under the previously tested conditions, namely 367 °C, 53.75 bar, an H_2 /wax ratio of 0.1275 kg/kg, and a WHSV of $2.5 \text{ kg}_{n-C, \text{in}} \text{ kg}_{\text{cat}}^{-1} \text{ h}^{-1}$, the ULiège inlet stream contains

98.09 mol% vapour, compared with 97.67 mol% for the Pellegrini inlet stream.




Figure 3.17: Comparison of the ULIège hydrocracking reactor inlet stream after olefin hydrogenation with the C_4 - C_{70} n-paraffin inlet distribution of Pellegrini et al. (2007) [4].

3.2.2 Kinetic profiles with ULIège inlet

With the kinetic model validated in Section 3.1 and the inlet stream characterised in Section 3.2.1, the model can now be applied to the ULIège case. The objective of this section is twofold: identify the operating point that maximises the outlet $C_8 - C_{16}$ mass fraction within the Pellegrini calibration domain, and characterise the behaviour of the reactor around that point. The discussion deliberately proceeds in three steps. The optimum itself is reported first (Section 3.2.2.1), followed by the kinetic profiles associated with it (Section 3.2.2.2) and by a sensitivity analysis that probes how robust this operating point is to perturbations of each of the four operating variables (Section 3.2.2.3). Throughout, the limitations inherited from the underlying model (the midpoint cracking stoichiometry and the partial-pressure assumption) are made explicit, because they condition the way the results must be read.

3.2.2.1 Optimal operating conditions

Applying the four-phase procedure of Section 2.4.2 to the full C_1 - C_{70+} ULIège inlet over the Pellegrini calibration domain, the iso-yield set collapses, after refinement and deduplication, to a single operating point:

$$T = 369.3 \text{ } ^\circ\text{C} \quad P = 60 \text{ bar} \quad \text{WHSV} = 3.00 \text{ kg}_{\text{n-C, in}} \text{ kg}_{\text{cat}}^{-1} \text{ h}^{-1} \quad \text{H}_2/\text{wax} = 0.150 \text{ kg/kg}$$


yielding 52.03 wt-% of $C_8 - C_{16}$ at the outlet and a C_{17+} mass conversion of $X_{C_{17+}} = 55.9 \%$. The inlet vapour mole fraction at these conditions is $\text{VFRAC} = 0.984$, comfortably above the 96 % threshold derived in Section 2.3.3, so the partial-pressure assumption underlying the kinetic model is admissible. Because the strict-optimality filter retains a single point, the economic criterion plays no role here and acts only as a tie-breaker for the broader cases reported in Section 3.3.

Two features of this result deserve attention before any further analysis. First, three of the four operating coordinates, namely pressure, WHSV, and the hydrogen-to-wax ratio, are pinned at the upper limit of the Pellegrini calibration domain, while the temperature, at 369.3 $^\circ\text{C}$ sits in the interior of the [343, 375] $^\circ\text{C}$ range. The result is therefore a corner solution: the search has converged to the boundary in three of the four directions, so that, within the response surface of the model, the yield would continue to increase if those bounds were relaxed (see Section 3.3). This interpretation calls for caution, however, since these bounds coincide with the limits of the experimental data on which the model was regressed; moving beyond them would amount to extrapolating an uncalibrated model rather than deriving a physically


meaningful recommendation. The interior temperature is the only genuine trade-off, and it admits a clear physical interpretation. The ULiège feed already contains 33.81 wt-% of $C_8 - C_{16}$ inherited from the upstream Fischer-Tropsch reactor. Raising the temperature accelerates the secondary cracking of this feed-borne kerosene faster than fresh kerosene is generated from the heavy ends and thereby depletes it, whereas lowering the temperature slows the cracking of the heavy ends from which additional kerosene is produced. The optimum is located at the temperature for which these two opposing effects balance. Since the position of this balance is governed by the kerosene content of the feed, the optimal temperature would shift with the inlet composition (see Section 3.3.4).

Second, the absolute yield reported here must be interpreted carefully. As emphasised in Section 3.1.4, the kinetic parameters were calibrated on the Pellegrini $C_4 - C_{70}$ feedstock and not refitted to the ULiège stream, which is lighter and contains a non-negligible $C_1 - C_4$ tail entirely absent from the Pellegrini calibration data. The figure of 52 wt-% should therefore be read as an order-of-magnitude indication of the achievable kerosene cut, not as a precise design value.

Figure 3.18 reports the inlet and outlet carbon-number distributions at this operating point. The feed peaks at C_{16} and decays smoothly above C_{20} , whereas the outlet shows a clear redistribution of mass from the heavy tail (C_{17+} falls from 63.23 to 27.71 wt-%) into the kerosene range, where the $C_8 - C_{16}$ lump grows from 33.8 to 52.0 wt-%. This redistribution is, however, not free: the model predicts 6.6 wt-% of $C_1 - C_4$ at the outlet, against essentially none at the inlet and an increase of the $C_5 - C_7$ naphtha fraction from 2.6 to 13.4 wt-%. The $C_1 - C_4$ yield represents a loss of carbon from the target product that would have to be either recycled or valorised as fuel gas; the $C_5 - C_7$ naphtha fraction can be routed to a separate aromatisation unit to be valorized, as discussed in Section 1.3.1. The increase in $C_8 - C_{16}$ of +18.2 wt-% is therefore the genuine contribution of the hydrocracking step itself, since the remaining 33.8 wt-% of the outlet kerosene fraction is inherited unchanged from the feed.



(a) Inlet and outlet carbon-number distribution on a mass basis.



(b) Lump breakdown of the inlet and outlet streams.

Figure 3.18: Inlet and outlet carbon-number distributions at the optimal operating point ($T = 369.3$ °C, $P = 60$ bar, $WHSV = 3.00$ $\text{kg}_{n-C, \text{in}} \text{kg}_{\text{cat}}^{-1} \text{h}^{-1}$, $H_2/\text{wax} = 0.150$).

3.2.2.2 Kinetic plots

Once the operating point is fixed, the integration of the material balance equations along the reactor produces the kinetic profiles of Figure 3.19.




Figure 3.19: Reactor profiles as a function of the space time $\tau = 1/\text{WHSV}$ at the recommended operating point.

The C_{17+} conversion (Figure 3.19a) displays an accelerating profile: the slope of the curve steepens with axial position even though the reactant is being depleted along the bed. This curvature is governed by the inhibition denominator ADSORB rather than by any increase in reactant partial pressure. The Langmuir adsorption constants $K_{L,n}$ and $K_{L,\text{iso}}$ grow exponentially with carbon number (Equations (3.1)–(3.2)), so the summation $\sum K_{L,n}p_n + \sum K_{L,\text{iso}}p_{\text{iso}}$ entering ADSORB is dominated by the heaviest paraffins. These same species carry the largest kinetic constants ($k_{\text{isom}} \propto n_c^{8.021}$, $k_{\text{cr}} \propto n_c^{5.851}$) and are therefore consumed preferentially in the first part of the reactor. As each heavy molecule is converted into two markedly less strongly adsorbed lighter fragments, ADSORB falls steeply along the bed. Since every reaction rate is inversely proportional to ADSORB (Equations (2.4)–(2.5)), this progressive de-inhibition raises the apparent rate of the remaining species faster than the depletion of the reactants lowers it, and the overall C_{17+} conversion accelerates with contact time.

The lump profiles (Figure 3.19b) confirm the mechanistic picture. The C_{17+} curve decreases monotonically, the $\text{C}_8 - \text{C}_{16}$ curve grows monotonically and crosses the C_{17+} curve near $\tau \approx 0.25 \text{ kg}_{\text{cat}} \text{h kg}_{\text{n-C, in}}^{-1}$, and the $\text{C}_5 - \text{C}_7$ and $\text{C}_1 - \text{C}_4$ curves grow slowly but visibly throughout at the end of the reactor. The latter two carry an important message: in this operating regime the kerosene production is not isolated, it is accompanied by a parallel light-gas make that grows monotonically with contact time. Pushing the conversion further by extending the bed length (i.e. lowering WHSV) would shift mass from C_{17+} towards both $\text{C}_8 - \text{C}_{16}$ and $\text{C}_1 - \text{C}_7$, with no mechanism in the model to favour one over the other; the kerosene selectivity for this case is therefore maximised at the maximum WHSV of the domain of calibration.

This boundary optimum is, however, a direct consequence of taking $\text{C}_8 - \text{C}_{16}$ selectivity as the priority objective: the procedure pushes towards the highest attainable $\text{C}_8 - \text{C}_{16}$ fraction regardless of what else is co-produced. Figure 3.19b nevertheless shows that around $\tau \approx 0.25 \text{ kg}_{\text{cat}} \text{h kg}_{\text{n-C, in}}^{-1}$, the C_{17+} fraction has already been largely consumed and has crossed below the kerosene curve, so most of the useful conversion is realised well before the edge of the calibration domain and, crucially, the sub- C_8 lumps ($\text{C}_1\text{--}\text{C}_4$ and $\text{C}_5\text{--}\text{C}_7$) remain marginal up to this point, their accumulation only accelerating further downstream. A natural perspective for future work would therefore be to extend the existing constraint set with an explicit upper bound on the light-gas fraction ($\text{C}_1\text{--}\text{C}_7$). Such an additional constraint would relocate the optimum to an interior space time, where the heavy fraction is substantially converted while light-gas production is still kept low.

Figure 3.20 compares the outlet distribution of total paraffins with the feed: the heavy tail of the wax (C_{22+}) is largely consumed and the distribution shifts toward the lighter $\text{C}_5\text{--}\text{C}_{20}$ range, confirming that the reactor performs the intended hydrocracking.




Figure 3.20: Outlet molar distribution of total paraffins ($n + \text{iso}$ combined), compared with the feed.

One feature of Figures 3.20–3.21 requires a critical comment, as it reveals a limitation inherited from the Pellegrini formulation that should not be hidden behind the yield.




Figure 3.21: Outlet molar distributions of n - and iso-paraffins at the recommended operating point.

The C_5 spike

The iso-paraffin distribution (Figure 3.21) exhibits a sharp, narrow peak at iso- C_5 , exceeding the kerosene-range concentrations by a factor of three. This spike is not a physical observation but a structural artefact of the cracking stoichiometry, and its origin is best understood as a kinetic bottleneck: iso- C_5 is fed rapidly from heavier species but consumed only slowly.

On the production side, iso- C_5 is fed by the central cracking of its heavier parents C_9 , C_{10} and C_{11} , abundant intermediates whose contributions appear explicitly as the production terms in the iso- C_5 material balance, Equation (2.16). On the consumption side, iso- C_5 is the lightest iso-paraffin allowed to crack, and it does so only through the special light-end rule of Reaction (2.9), $\text{iso-}C_5 + H_2 \rightarrow 0.5 C_1 + 0.5 C_2 + 0.5 C_3 + 0.5 \text{iso-}C_4$. Because k_{cr} increases steeply with carbon number, $k_{cr}(5)$ is the smallest cracking constant of the entire series, so iso- C_5 is removed far more slowly than it is generated from C_9 – C_{11} . This imbalance between a fast inflow and a slow outflow makes iso- C_5 a point of mass accumulation, which is the direct cause of the spike. The accumulation is then mirrored on the n -paraffin side: the hydroisomerization step $n\text{-}C_5 \rightleftharpoons \text{iso-}C_5$ is reversible and close to equilibrium, so any local excess of iso- C_5 drives the reverse reaction and inflates $n\text{-}C_5$. The $n\text{-}C_5$ peak is therefore an equilibrium reflection of the iso- C_5 bottleneck rather than an independent feature.

The narrower peaks at C_3 and C_4 share this origin but behave as pure dead ends: C_1 , C_2 and C_3 are terminal species that never crack, while iso- C_4 does not crack and $n\text{-}C_4$ only isomerises to iso- C_4 . The

graphical convention adopted by the authors, which consists in grouping iso-C₄ with n-C₄, is retained here. This explains why the iso-C₄ fraction is equal to zero. None of these light species has any cracking exit, so they accumulate monotonically from every cracking event that produces them; n-C₃, for instance, receives mass from the central cracking of C₆ (two C₃) and C₇ (one C₃) on top of the 0.5 C₃ delivered by the iso-C₅ rule, and is never depleted. These light-end peaks are thus a structural feature of the midpoint cracking stoichiometry combined with the special iso-C₅ rule, not a reflection of the underlying chemistry. They would be smoothed out by the breakage probability function of Gamba et al. [6] (Section 1.4.2.2.2), and are reported here as the model produces them, without being read as quantitative predictions of the C₃, C₄ and C₅ yields.

3.2.2.3 Sensitivity analysis

This section probes the robustness of the recommended optimum to perturbations of the four operating variables, with two distinct objectives. The two-dimensional slices map the response surface of the kerosene yield in the neighbourhood of the optimum and check that the operating point lies in a smooth region rather than on a numerical artefact. The one-dimensional sweeps, by contrast, extend each variable to the bounds of the broader literature domain and expose the qualitative behaviour of the model away from its calibration range, where threshold effects emerge that have direct consequences for reactor design.

3.2.2.3.1 Two-dimensional sensitivity slices

Figure 3.22 reports the six two-dimensional slices of the C₈ – C₁₆ outlet mass fraction obtained by varying each pair of operating variables while holding the other two fixed at the recommended values.




Figure 3.22: Two-dimensional sensitivity slices of the C₈ – C₁₆ outlet mass fraction around the recommended operating point ($T = 369.3$ °C, $P = 60$ bar, $\text{WHSV} = 3.00 \text{ kg}_{n-C,\text{in}} \text{ kg}_{\text{cat}}^{-1} \text{ h}^{-1}$, $\text{H}_2/\text{wax} = 0.150 \text{ kg/kg}$).

Across the six panels the slices share a common structure, which is worth setting out before the individual axes are examined. The colour encodes the C₈ – C₁₆ outlet mass fraction, the white lines are its iso-yield contours, and the recommended optimum is marked in each panel. Every slice separates into two

regions: a productive region, in which the outlet fraction rises to a 45–52 wt-% plateau, and a near-zero region, in which the kerosene has been overcracked to light gas. The iso-yield contours are widely spaced within the productive region and crowd into a narrow band at its edge, so the two regions are divided by a sharp boundary rather than through a gradual transition.

This sharpness has a kinetic origin. The $C_8 - C_{16}$ cut is an intermediate of the cracking sequence (C_{17+} heavy ends $\rightarrow$ kerosene $\rightarrow$ light gas): its outlet level results from the balance between the rate at which it is formed from the heavy ends and the rate at which it is itself cracked to lighter products. In the LHHW rate law, every cracking rate is lowered by the species adsorbed on the catalyst, an inhibition dominated by the heaviest paraffins, which adsorb most strongly (largest Langmuir constants). As these heavy paraffins are consumed, the inhibition they exert weakens and all the cracking rates increase, so the reaction accelerates as it proceeds (Section 3.2.2.3.2). This self-acceleration produces a threshold behaviour. Below a certain severity, the conversion stays moderate and the kerosene is preserved; once that severity is exceeded, the C_{17+} conversion runs close to completion, the kerosene is overcracked to light gas, and the outlet $C_8 - C_{16}$ drops abruptly. The high-yield plateau and the overcracked region are thus separated by a sharp boundary, with the optimum on the plateau side of it.

The position of this boundary is far more sensitive to some of the four operating variables than to others, as the three observations below quantify.

Temperature is the dominant axis of activity

In each of the three panels involving T ($T \times P$, $T \times \text{WHSV}$, $T \times H_2/\text{wax}$), a sharp transition from the high-yield plateau to zero yield is visible in the temperature direction, located between 355 and 370 °C. The recommended optimum at $T = 369.3$ °C sits at the edge of an active band, well within the plateau but close enough to its side that a temperature increase of 5–10 °C would shift the reactor into the overcracking (0 wt-%) regime. The implications for reactor control are discussed in Section 3.2.2.3.3 together with the corresponding one-dimensional sweeps.

Pressure and WHSV are not inactive

In the three panels that do not involve T ($\text{WHSV} \times P$, $P \times H_2/\text{wax}$, $\text{WHSV} \times H_2/\text{wax}$), with the temperature held at 369.3 °C, the response is far from flat: the $C_8 - C_{16}$ fraction varies by approximately 50 wt-% across the $\text{WHSV} \times P$ panel, with a clear threshold near $P \approx 50\text{--}55$ bar and $\text{WHSV} \approx 2.6 \text{ kg}_{\text{n-C, in}} \text{ kg}_{\text{cat}}^{-1} \text{ h}^{-1}$ below which the yield collapses. The interpretation is the same as the temperature transition but driven by different mechanisms: at low pressure, p_{H_2} in the ADSORB denominator falls, the rate accelerates, and the formed kerosene is overcracked; at low WHSV, the extended contact time produces the same overcracking outcome. The recommended pressure (60 bar) and WHSV ($3.0 \text{ kg}_{\text{n-C, in}} \text{ kg}_{\text{cat}}^{-1} \text{ h}^{-1}$) both sit at the upper bound of the Pellegrini calibration domain, on the high-yield plateau and as far as possible from the overcracking transition.

The H_2/wax ratio is essentially inactive

The three panels involving H_2/wax show that, of the four operating variables, this ratio has the least influence on the $C_8 - C_{16}$ outlet mass fraction. In the $T \times H_2/\text{wax}$ panel the iso-yield contours are essentially vertical, so the outlet fraction is set by temperature and barely changes along the ratio axis. In the $P \times H_2/\text{wax}$ and $\text{WHSV} \times H_2/\text{wax}$ panels the ratio is likewise inactive across both the high-yield plateau and the dead zone, and acquires leverage only within the narrow overcracking transition, where a higher ratio displaces the boundary slightly towards lower pressure (respectively lower WHSV). Since the recommended optimum lies well inside this plateau, the residual sensitivity remains limited. The one-dimensional sweep reported in Figure 3.27 reports that doubling the H_2/wax ratio increases the $C_8 - C_{16}$ outlet mass fraction by only about 4 wt-%, while the other operating variables are kept fixed. This

behaviour is consistent with the partial-pressure assumption of the model: in the absence of a fugacity-based treatment, the H_2 /wax ratio has no thermodynamic lever on the heavy-end partitioning between the two phases, and acts only through the hydrogen partial pressure in the ADSORB denominator, which is precisely why its influence is confined to the pressure- and contact-time-controlled transition. This near-inactivity is therefore not a property of the underlying chemistry but a known limitation of the present implementation: the fugacity-based model of Pellegrini et al. (2008) [4] predicts a stronger dependence.

Beyond the influence of the individual operating variables, the structure of the optimum is itself informative. The strict search returns a single grid point in the high-yield region. This does not imply a flat response surface there. Were the surface genuinely flat near the maximum, the strict tolerance ($\Delta_{iso} = 0.05$ wt-%) would already admit several tied points; that it admits only one indicates a sharp maximum rather than a plateau. The recommended optimum, moreover, sits at a corner of the domain along three of the four axes (Section 3.2.2.1): along these three axes no grid point lies beyond the optimum, and along the remaining direction none falls within 0.05 wt-% of it. The isolation of the strict-scenario optimum therefore reflects its corner location combined with the deliberately narrow tolerance, not any flatness of the surface. The relaxed scenario of Section 3.3, which widens the tolerance to 5 wt-%, immediately returns several hundred grid points, consistent with a response that descends appreciably away from the maximum across the high-yield region.

3.2.2.3.2 Reactor profiles under varying operating conditions

The two-dimensional slices map only the outlet state of the reactor. A complementary view is obtained by integrating the C_{17+} mass conversion along the reactor for several values of each operating variable in turn, the three others being held at the recommended optimum.

Temperature (Figure 3.23a)

The conversion profiles progressively diverge as the temperature increases. This spread reflects the Arrhenius dependence of the kinetic constants k_{iso} and k_{cr} : temperature scales every rate constant through the factor $\exp(-E/RT)$, so a higher temperature lifts the whole profile towards higher conversion. Three regimes can be distinguished. Below ~ 345 °C, the rate constants are too small to produce any measurable conversion within the residence time $\tau_{out} = 0.333 \text{ kg}_{n-C_{10},in}^{-1} \text{ kg}_{cat} \text{ h}$ allowed by the recommended WHSV, and the profiles stay near zero throughout the reactor. Between 345 and 365 °C a productive regime opens: each profile develops a clear upward curvature, accelerating along the bed. This curvature is not an Arrhenius effect, temperature being constant along an isothermal reactor, but a de-inhibition effect: the heaviest paraffins, which carry both the largest kinetic constants and the largest Langmuir constants, are consumed preferentially in the first part of the bed, and their conversion into weakly adsorbed lighter fragments lowers the ADSORB denominator, which raises the apparent rate of every remaining species with axial position. Above ~ 370 °C, $X_{C_{17+}}$ exceeds 80 wt-% well before the reactor outlet, signalling that the operating point would shift to the overcracking regime if the temperature were allowed to drift upward.

Pressure (Figure 3.23b)





The profiles form a descending family with increasing pressure, exactly as the LHHW formulation predicts: p_{H_2} enters the ADSORB denominator linearly, so a higher total pressure increases the denominator and reduces every reaction rate uniformly. At $P = 35$ bar, the reactor is in the overcracking regime ($X_{C_{17+}} > 0.8$ at the outlet); at $P = 80$ bar the conversion remains modest throughout. The recommended pressure (60 bar, upper bound of the Pellegrini domain) sits on the upper side of this family, providing the kinetic protection of the formed kerosene against secondary cracking. This protective mechanism is the kinetic reason behind the corner location of the optimum in pressure.

H₂/wax ratio (Figure 3.23c)

Across the four profiles, the spread is far smaller than in the temperature and pressure sweeps (panels 3.23a and 3.23b): raising the hydrogen-to-wax ratio from 0.06 to 0.15 kg/kg changes the outlet conversion only slightly. This confirms the conclusion already drawn from the two-dimensional slices: under the partial-pressure assumption, the H₂/wax ratio influences the reactor only through its effect on p_{H_2} in the inhibition denominator, with no separate effect on the heavy-end partitioning.

WHSV (Figure 3.23d)

The WHSV sweep is structurally different from the three others: WHSV does not enter the right-hand side of the material balance equations (Section 2.3.1), it only fixes the upper integration limit $\tau_{out} = 1/\text{WHSV}$. Consequently, the four WHSV cases lie on a single common profile at $T = 369.3$ °C, $P = 60$ bar and $H_2/\text{wax} = 0.150$ kg/kg, cut at different final times. The figure therefore reports the profile with the four operating points (τ_{out}, X_{out}) explicitly marked, rather than four redundant curves. Reading the figure in this way makes the WHSV trade-off explicit: lowering WHSV moves the operating point rightward along the curve, which increases $X_{C_{17+}}$ but, as established by the corresponding lump sweep below, at the price of overcracking the formed kerosene to light gases. The recommended WHSV = $3.0 \text{ kg}_{n-C, in} \text{ kg}_{cat}^{-1} \text{ h}^{-1}$ corresponds to the latest point along the profile before this trade-off turns unfavourable.

(a) Temperature sweep, $T = 305\text{--}375$ °C.(b) Pressure sweep, $P = 35\text{--}80$ bar.(c) H₂/wax sweep, 0.06–0.15 kg/kg.

(d) WHSV cut points on the common profile.

Figure 3.23: Mass conversion of the C₁₇₊ fraction as a function of the space time $\tau = 1/\text{WHSV}$, for variations of each operating variable around the recommended point. In every panel, the three variables not being swept are held at $T = 369.3$ °C, $P = 60$ bar, WHSV = $3.00 \text{ kg}_{n-C, in} \text{ kg}_{cat}^{-1} \text{ h}^{-1}$, $H_2/\text{wax} = 0.150$ kg/kg.

3.2.2.3.3 One-dimensional sweeps

The two-dimensional slices reveal the shape of the response surface near the optimum. To make the threshold behaviour explicit, each of the four variables is now swept individually across its full literature range while the other three are held at the recommended optimum. The outlet composition is reported as four lumps ($C_1 - C_4$, $C_5 - C_7$, $C_8 - C_{16}$, C_{17+}), together with the inlet values (dotted horizontal lines) for direct comparison.

The temperature sweep (Figure 3.24) is the most informative of the four. Three regimes can be distinguished. Below approximately 345 °C the outlet lumps coincide with the inlet values: the reaction rates are too small to produce any measurable conversion over the residence time $\tau = 0.33 \text{ kg}_{\text{n-C},\text{in}}^{-1} \text{ kg}_{\text{cat}} \text{ h}$ allowed by $\text{WHSV} = 3.00 \text{ kg}_{\text{n-C},\text{in}} \text{ kg}_{\text{cat}}^{-1} \text{ h}^{-1}$. It can therefore be observed that extrapolating the temperature to the range reported in the literature, i.e., below 343 °C, seems to yield a consistent behaviour in this case. Between 345 and 370 °C a productive window opens: C_{17+} falls monotonically, $C_8 - C_{16}$ rises monotonically and reaches its maximum near 369–370 °C, while $C_1 - C_4$ and $C_5 - C_7$ grow at a controlled pace. Above ~ 372 °C the kerosene yield collapses while the $C_1 - C_4$ lump rises rapidly to dominate the outlet, signalling that the same kinetic constants that enable cracking of the heavy ends now also drive secondary cracking of the already-formed kerosene. The productive window between the two regimes is narrow, approximately 20 °C wide for this feed, and the optimum identified in Section 3.2.2.1 sits within 3 °C of its upper edge.

This last point should not be underestimated. In an isothermal plug-flow reactor with strongly Arrhenius-activated kinetics, an upward temperature deviation of only 3 °C from the optimal point, which can readily occur in practice, would be sufficient to shift the system onto the overcracking branch and redirect the kerosene mass toward the $C_1 - C_4$ lump. The physical origin of this sharp transition cannot be isolated from the influence of the simplified midpoint stoichiometry within the present modelling framework; experimental hydrocracking campaigns typically report smoother yield curves [45], which suggests that the breakage probability function and the full LHHW rate expression of Gambaro et al. [7] would broaden this transition. Even allowing for this, the qualitative message of Figure 3.24 remains: the recommended operating point requires tight thermal control, which has direct implications for reactor design.




Figure 3.24: Sweep of the temperature, $P = 60$ bar, $\text{WHSV} = 3.00 \text{ kg}_{\text{n-C},\text{in}} \text{ kg}_{\text{cat}}^{-1} \text{ h}^{-1}$, $\text{H}_2/\text{wax} = 0.150 \text{ kg/kg}$.

The pressure sweep (Figure 3.25) is qualitatively the mirror image of the temperature sweep, but driven by a different mechanism. Below approximately 52 bar the reactor runs in the overcracking regime, with $C_1 - C_4$ approaching 100 wt-% of the outlet; above ~ 55 bar the kerosene yield rises and levels off/decreases very slightly near 50 wt-% for $P \geq 65$ bar. This trend is the expected consequence of the hydrogen partial pressure in the ADSORB denominator: as total pressure increases, p_{H_2} rises, the denominator grows, and the overall rate falls, which protects the formed kerosene from secondary cracking.

The recommended pressure of 60 bar sits at the upper bound of the Pellegrini calibration domain and lies just inside the protective regime. Extrapolation to the broader literature domain, in which a coherent behaviour is observed, indicates that the slight yield decrease extends beyond 60 bar. This can be attributed to the fact that, above a certain pressure, increasing the total pressure raises the hydrogen partial pressure appearing in the ADSORB denominator, thereby adversely affecting the dehydrogenation equilibrium and reducing the concentration of reactive olefinic intermediates at the acid sites, as discussed in Section 1.4.4.




Figure 3.25: Sweep of the total pressure, $T = 369.3\text{ }^{\circ}\text{C}$, $\text{WHSV} = 3.00\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$, $\text{H}_2/\text{wax} = 0.150\text{ kg/kg}$.

The WHSV sweep (Figure 3.26) shows that below $\text{WHSV} \approx 2.6\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$ the extended contact time pushes the system into the overcracking branch, C₁ – C₄ rises sharply, and C₈ – C₁₆ falls below its inlet value: more contact time degrades the yield rather than enhancing it. Between $\text{WHSV} = 2.7$ and $4\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$, the kerosene yield is relatively constant at $\sim 50\text{ wt-\%}$. The recommended value $\text{WHSV} = 3.0\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$ therefore lies on the extremity of the Pellegrini et al. calibration domain. This boundary location is consistent with the model response: decreasing the WHSV would increase the contact time, promote overcracking, and would not improve the kerosene yield. Within the Pellegrini et al. calibration domain, selecting $\text{WHSV} = 3.0\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$ therefore minimises the required catalyst mass without reducing the predicted yield.

Once again, extrapolation from the Pellegrini et al. calibration domain to the broader literature domain provides a coherent result. As the WHSV increases, the contact time becomes shorter, progressively limiting conversion until the mass fraction approaches its initial value.




Figure 3.26: Sweep of the WHSV, $T = 369.3$ °C, $P = 60$ bar, $H_2/\text{wax} = 0.150$ kg/kg.

The H_2/wax sweep (Figure 3.27) confirms, finally, the conclusion already drawn from the two-dimensional slices. Across the range $H_2/\text{wax} = 0.06$ to 0.15 kg/kg, the four lumps vary by a few percentage points: the $C_8 - C_{16}$ fraction rises by approximately +4 wt-% when the ratio is doubled, while $C_5 - C_7$ decreases by a similar amount. This near-flat response is the structural signature of the partial-pressure assumption: as explained previously, in this model, the hydrogen ratio influences the rates only through the inhibition denominator, with no thermodynamic action on the vapour-liquid partition of the heavy paraffins. In a fugacity-based model accounting for VLE, as in Pellegrini et al. (2008) [4], the same sweep would produce a stronger dependence, because increasing H_2/wax would strip heavy paraffins towards the vapour phase, increase their effective fugacity at the catalyst surface, and accelerate their conversion. The model's insensitivity to H_2/wax should be therefore regarded as a known limitation rather than as a physical prediction.





Figure 3.27: Sweep of the hydrogen-to-wax mass ratio, $T = 369.3$ °C, $P = 60$ bar, $\text{WHSV} = 3.00$ kg_{n-C,in} kg_{cat}⁻¹ h⁻¹.

3.2.2.4 Thermodynamic admissibility under perturbations

The kinetic model is only valid under the partial-pressure approximation introduced in Section 2.3.3, which requires the overall vapour mole fraction of the mixture to remain above 96 %. A reactor operated in practice will never sit exactly at the nominal operating point, so this admissibility condition must be verified not only at the recommended point itself, but also across the perturbation range over which the kinetic results have been reported. To this end, the inlet flash was re-evaluated in Aspen Plus with the SRK equation of state along the temperature and pressure axes of the literature domain, with $H_2/\text{wax} = 0.150$ held at its nominal value.

Two conclusions can be drawn from Figure 3.28. First, across the whole pressure range [35, 80] bar, the vapour fraction varies between approximately 0.978 and 0.992 and never approaches the 0.96 threshold; the kinetic model is therefore admissible everywhere within the pressure sweep of Figure 3.25. Second, the temperature sweep also remains above this threshold over the whole range, leading to the same conclusion for Figure 3.24: the model remains valid throughout the investigated domain. The recommended hydrogen ratio $H_2/wax = 0.150$ is the dominant factor: with 94.65 mol-% of the inlet being hydrogen, the heavy hydrocarbon liquid fraction is small enough that the global VFRAC stays comfortably above the admissibility line over the entire operating window. Lowering H_2/wax would erode this margin and is the reason why some scenario optima of Section 3.3 that minimise hydrogen consumption fail the thermodynamic check.



(a) Vapour fraction vs total pressure at $T = 369.3$ °C.

(b) Vapour fraction vs temperature at $P = 60$ bar.

Figure 3.28: Inlet overall vapour mole fraction (VFRAC) and liquid mole fraction (LFRAC) computed in Aspen Plus with the SRK equation of state, at $H_2/wax = 0.150$.

3.2.3 Reactor sizing

The reactor sizing methodology established in Section 2.5 is now applied to the recommended operating point identified in Section 3.2.2.1 for the full C_1 – C_{70+} ULiège feed:

$$T = 369.3 \text{ °C} \quad P = 60 \text{ bar} \quad WHSV = 3.00 \text{ kg}_{n-C_{10},in} \text{ kg}_{cat}^{-1} \text{ h}^{-1} \quad H_2/wax = 0.150 \text{ kg kg}^{-1}$$

Before reporting any reactor dimension, it should be emphasised that the present sizing is based on a fixed feed, corresponding to the inlet hydrocarbon stream defined in Section 3.2.1, with a mass flow rate of

$$\dot{m}_{wax,in} = 0.113918 \text{ kg h}^{-1} \quad (3.15)$$

which is the pilot-scale flow of the ULiège installation.

3.2.3.1 Catalyst mass and hydrogen co-feed

Applying Eq. (2.40), the catalyst mass required to operate at the recommended WHSV is

$$m_{cat} = \frac{\dot{m}_{wax,in}}{WHSV} = \frac{0.113918}{3.00} \approx 37.97 \text{ g} \quad (3.16)$$

The hydrogen co-feed follows from the hydrogen-to-wax mass ratio through Eq. (2.46):

$$\dot{m}_{H_2} = \left(\frac{H_2}{wax} \right) \dot{m}_{wax,in} = 0.150 \times 0.113918 \approx 17.09 \text{ g h}^{-1} \quad (3.17)$$

As established in the introduction of Section 2.5, hydrogen does not enter the geometric closure of the reactor: it percolates through the inter-particle voids already accounted for in the bed volume. It is therefore retained here as an operating specification rather than as a geometric sizing variable, even though it remains essential to the LHHW rate expressions and to the gas–liquid hydrodynamics of trickle-bed operation.

3.2.3.2 Bulk density of the catalytic bed

Converting the catalyst mass into a packed-bed volume requires the bulk density of the bed, a quantity seldom reported in the literature. It was therefore estimated from the reference bed of Pellegrini et al. [4], of known catalyst mass and geometry. The catalyst is the bifunctional system described in Section 1.4.1: platinum (0.6 wt%) supported on amorphous silica–alumina, employed as a fine powder of particle diameter

$$d_p = 0.625 \text{ mm} \quad (3.18)$$

as reported for the experimental campaigns of the Pellegrini series. The reference bed, of internal diameter $D_{\text{ref}} = 16 \text{ mm}$ and bed height $H_{\text{ref}} = 86 \text{ mm}$, occupies the cylindrical volume

$$V_{\text{bed,ref}} = \frac{\pi D_{\text{ref}}^2 H_{\text{ref}}}{4} = \frac{\pi (1.6)^2 (8.6)}{4} \approx 17.3 \text{ cm}^3 \quad (3.19)$$

and contains $m_{\text{cat,ref}} = 9 \text{ g}$ of catalyst, which yields a bulk density

$$\rho_{\text{bulk}} = \frac{m_{\text{cat,ref}}}{V_{\text{bed,ref}}} = \frac{9}{17.3} \approx 520.5 \text{ kg m}^{-3} \quad (3.20)$$

This value is physically consistent with a randomly packed bed of porous particles [92]. The bulk density reflects both the density of the catalyst particles ρ_p and the inter-particle void fraction ε_b of the packing, the two being related by $\rho_{\text{bulk}} = (1 - \varepsilon_b) \rho_p$. Assuming a representative void fraction $\varepsilon_b = 0.40$, in line with the correlation of Dixon [95], the particle density is recovered as

$$\rho_p = \frac{\rho_{\text{bulk}}}{1 - \varepsilon_b} \approx 867.5 \text{ kg m}^{-3} \quad (3.21)$$

This figure lies far below the density of dense alumina ($3900\text{--}4000 \text{ kg m}^{-3}$ [96]), as expected from the extensive internal mesoporosity of the silica-alumina support [92]. It is reported here only as a physical-consistency check.

3.2.3.3 Packed-bed volume and reactor geometry.

With the bulk density known, the packed-bed volume follows directly from Eq. (2.44):

$$V_R = \frac{m_{\text{cat}}}{\rho_{\text{bulk}}} = \frac{37.97}{0.520} \approx 73.0 \text{ cm}^3 \quad (3.22)$$

The reactor is modelled as a cylindrical packed bed, whose volume relates the internal diameter D and the bed height H through Eq. (2.45), $V_R = \pi D^2 H / 4$. Since V_R is already fixed by the catalyst mass, this relation leaves a single geometric degree of freedom. Following the procedure of Section 2.5.3, the aspect ratio is set to its minimum admissible value $H/D = 5$, which corresponds to the shortest bed for which axial dispersion and wall channelling remain negligible, and hence to the geometry of lowest pressure drop among those compatible with plug-flow behaviour. Combining $V_R = \pi D^2 H / 4$ with $H = 5D$ yields

$$D = \left(\frac{4V_R}{5\pi} \right)^{1/3} \approx 26.5 \text{ mm} \quad (3.23)$$

$$H = 5D \approx 132.4 \text{ mm} \quad (3.24)$$

The resulting tube-to-particle diameter ratio is

$$\frac{D}{d_p} = \frac{26.5}{0.625} \approx 42.4 \quad (3.25)$$

comfortably above the lower bound $D/d_p \geq 25$ imposed in Section 2.5.3 to limit wall channelling. Both hydrodynamic criteria, $D/d_p \geq 25$ and $H/D \geq 5$, are therefore satisfied.

3.2.3.4 Synthesis and interpretation

These reported results are calculated in such a way that the resulting reactor sizing is governed almost entirely by the selected WHSV: through Eq. (2.40) it fixes the catalyst mass and, in turn, the packed-bed volume, whereas the hydrogen flow, although a central operating variable, does not enter the geometric closure. For the recommended operating point applied to the full C_1 - C_{70+} ULiège feed, the hydrocracking reactor corresponds to a cylindrical packed bed of approximately 26.5 mm internal diameter and 132 mm bed height, holding close to 38 g of catalyst in a catalytic volume of about 73 cm³, as shown in Figure 3.29.

These dimensions should nonetheless be interpreted with caution. They were obtained for the sole purpose of providing an order-of-magnitude estimate of the reactor size, and several simplifications underlie them. The bed volume rests on a bulk density transposed from a single reference bed; the geometry is derived from idealised plug-flow criteria rather than from a detailed hydrodynamic analysis; and, as emphasised in Section 2.5.3, a construction-grade design would require, among others, an explicit pressure-drop evaluation through the Ergun equation and a mechanical sizing of the vessel at the operating pressure. Moreover, since the catalyst mass scales linearly with the feed throughput, the figures reported above are specific to the pilot-scale flow of 0.114 kg h⁻¹; a unit sized for the industrial NKL target would be larger by several orders of magnitude. The present geometry should therefore be read as an indicative result, consistent with the modelling assumptions of this work, and not as a finalised reactor design.




Figure 3.29: Schematic representation of the cylindrical packed-bed reactor at the recommended operating point (full C_1 - C_{70+} ULiège feed, $T = 369.3$ °C, $P = 60$ bar, $WHSV = 3.00$ kg_{n-C,in} kg_{cat}⁻¹ h⁻¹, $H_2/wax = 0.15$ kg/kg).

3.3 Optimal operating conditions with other considerations

This section reports the operating conditions that maximise the C_8 - C_{16} outlet mass fraction across the eight cases defined in Section 2.4.1: two feed streams (the full C_1 - C_{70+} ULiège stream and solely its C_{17+} fraction), two exploration domains (the Pellegrini calibration domain and the extended literature domain) and two precision scenarios (strict and relaxed optimality). The four-phase procedure of Section 2.4.2 is applied independently to each case, with the primary objective being the C_8 - C_{16} mass fraction at the reactor outlet.

Two practical aspects of the procedure must be recalled before the results are examined.

First, the thermodynamic admissibility condition introduced in Section 2.3.3 (overall vapour mole fraction $VFRAC \geq 0.96$) is not a constraint enforced inside the search itself. Since the admissibility criterion is evaluated through an Aspen Plus flash calculation rather than enforced directly during the Python optimisation, thermodynamic admissibility is assessed in a subsequent feasibility-screening step applied to the economically ranked set of refined optima. The recommended operating point is then selected as the lowest-cost refined candidate satisfying this criterion. As a direct consequence, when the most economical candidate fails the thermodynamic test, it is rejected and the next-ranked candidate is examined.

Second, the central limitation already stated in Section 2.4.1 and in the Introduction of this chapter must be kept in mind throughout the discussion. The kinetic parameters were calibrated on the Pellegrini et al. (2007) feedstock and were not refitted to the ULiège stream. The operating conditions and yields reported below therefore quantify the behaviour of the model rather than a guaranteed plant performance, and the value of this section lies as much in the trends it exposes as in the absolute figures it produces. The extent of this limitation is examined in Section 3.3.4, where the search is repeated on the Pellegrini calibration feed to observe the influence of the feed composition on the results.

3.3.1 Strict-optimality scenario

In the strict scenario, the tolerances are set to $\Delta_{\text{iso}} = 0.05$ wt-% and $\Delta_{\text{best}} = 0.01$ wt-%, so that only operating points indistinguishable from the global maximum within the numerical precision of the simulation are retained. The procedure is therefore expected to return either a single best point or a small set of genuinely equivalent points among which the economic criterion acts as a tie-breaker. The refined optima obtained for the four feed/domain combinations are reported in Table 3.1.

Table 3.1: Refined optima retained under the strict-optimality scenario.

Feed	Optimum No.	T [°C]	P [bar]	H ₂ /wax [kg/kg]	WHSV [kg _{n-C₁₀} kg _{cat} ⁻¹ h ⁻¹]	VFRAC	C ₈ -C ₁₆ %wt	Net C ₈ -C ₁₆ %wt	X _{C₁₇} ⁺ %wt	Optimum validity
Feed C₁-C₇₀₊										
Pell	1	369.3	60	0.15	3	0.984	52.033	18.220	55.9	Valid
Litt	1	375	68.8	0.06	4	0.929	52.539	18.726	56.2	Not valid
	2	375	80	0.103	3.23	0.963	52.539	18.726	56.2	Valid
	3	375	63	0.142	4	0.984	52.539	18.726	56.2	Valid
	4	375	69.1	0.15	3.64	0.983	52.538	18.725	56.2	Valid
Feed C₁₇₊										
Pell	1	375	53.4	0.06	3	0.918	38.857		67.5	Not valid
	2	375	50.1	0.115	3	0.967	38.857		67.5	Valid
	3	375	53.8	0.15	2.75	0.977	38.858		67.5	Valid
	4	375	60	0.15	2.46	0.974	38.858		67.5	Valid
Litt	1	375	43	0.139	3.45	0.978	38.858		67.5	Valid
	2	375	43.7	0.15	3.38	0.981	38.858		67.5	Valid
	3	375	40.7	0.15	3.63	0.982	38.858		67.5	Valid
	4	375	49.9	0.15	2.96	0.978	38.858		67.5	Valid

Within the Pellegrini domain applied to the full feed, the procedure returns a single refined optimum, already analysed in Section 3.2.2.1. For the three other combinations, the iso-yield set contains several refined optima that share the same yield within the numerical precision of the integration (visible in Table 3.1 from the identical C₈-C₁₆ values across candidates), and differ essentially through their hydrogen ratio and pressure. This is the flat-response situation anticipated in Section 2.4.2: once the dominant variables (T at 375 °C, WHSV near its upper bound) are fixed, the response surface flattens and the remaining variables form a plateau within which the economic criterion becomes the only discriminant.

Within this plateau, the procedure ranks candidates by ascending hydrogen ratio. In two of the four combinations, the cheapest candidate fails the post-hoc VFRAC check.

A second feature visible in the table is the systematic location of the optima on the upper temperature bound of each domain. In every case examined except the Pellegrini C₁-C₇₀₊ combination, the refined temperature sits exactly at 375 °C. This systematic boundary location indicates that, within the model, the kinetic gain from operating at higher temperatures is not offset by a comparable loss from secondary cracking until the upper bound is reached. Two distinct mechanisms account for this. The first applies to the extended domain: by raising the WHSV ceiling from 3 to 4 kg_{n-C₁₀} kg_{cat}⁻¹ h⁻¹, the extended domain allows the search to pair the highest temperature with a higher space velocity. The shorter contact time that results leaves the freshly formed kerosene less exposed to consecutive cracking, so the temperature can be driven to its upper bound without incurring the overcracking penalty that would otherwise cap it.

The second mechanism applies to the C_{17+} feed on both domains: since this feed is, by construction, free of C_8 – C_{16} at the inlet, no feed-borne kerosene is present that a high temperature could overcrack, and the optimisation is free to maximise the kerosene generated within the reactor at the temperature ceiling.

3.3.2 Relaxed-optimality scenario

In the relaxed scenario, the iso-yield threshold is widened to $\Delta_{iso} = 5$ wt-% and the post-refinement filter to $\Delta_{best} = 3$ wt-%, so that operating points yielding a few weight-percent below the maximum are deliberately admitted. The purpose of this scenario is to test whether a modest, accepted concession on yield opens access to materially more economical operating conditions. The iso-yield sets accordingly grow by orders of magnitude with respect to the strict case, containing 669, 274, 295 and 107 grid points for the four feed/domain combinations, against a handful in the strict scenario. The refined optima retained after farthest-point sampling, local refinement and ranking are reported in Table 3.2.

Table 3.2: Refined optima retained under the relaxed-optimality scenario.

Feed	Optimum No.	T [°C]	P [bar]	H ₂ /wax [kg/kg]	WHSV [kg _{n-C.in} kg _{cat} ⁻¹ h ⁻¹]	VFRAC	C ₈ –C ₁₆ %wt	Net C ₈ –C ₁₆ %wt	X _{C₁₇} ⁺ %wt	Optimum validity
Feed C₁–C₇₀₊										
Pell.	1	367.7	60	0.06	3	0.932	51.891	18.078	55.6	Not valid
	2	360	35	0.15	3	0.990	51.205	17.392	54.5	Valid
	3	369.3	60	0.15	3	0.984	52.033	18.220	55.9	Valid
	4	350.6	60	0.15	1	0.979	50.338	16.525	53.0	Valid
Litt.	1	375	68.8	0.06	4	0.929	52.539	18.726	56.2	Not valid
	2	375	80	0.06	3.44	0.920	52.539	18.726	56.2	Not valid
	3	375	80	0.142	3.15	0.978	52.539	18.726	56.2	Valid
	4	355.4	80	0.15	1	0.975	50.784	16.971	53.8	Valid
Feed C₁₇₊										
Pell.	1	375	57.3	0.06	2.80	0.916	38.857		67.5	Not valid
	2	375	60	0.133	2.48	0.969	38.857		67.5	Valid
	3	359.8	60	0.15	1	0.971	36.868		66.2	Valid
	4	375	60	0.15	2.46	0.974	38.858		67.5	Valid
Litt.	1	363.2	80	0.06	1	0.900	37.325		66.6	Not valid
	2	375	40.1	0.06	4	0.929	38.857		67.5	Not valid
	3	375	37	0.15	4	0.984	38.858		67.5	Valid
	4	375	80	0.15	1.85	0.968	38.857		67.5	Valid
	5	364.6	80	0.15	1	0.966	37.507		66.8	Valid

Three observations follow from Table 3.2.

The widening of the tolerance has only a marginal effect on the recommended yield. In three of the four cases the relaxed recommendation returns exactly the same outlet C_8 – C_{16} fraction as the strict scenario; in the fourth, the C_1 – C_{70+} feed on the Pellegrini domain, it costs 0.8 wt-%, a case examined in detail in the fourth observation of Section 3.3.3. The effect is most pronounced for the C_{17+} feed, where the strict and relaxed recommendations on both domains return an essentially identical yield of 38.86 wt-%.

A second observation is that, even when the recommended yield is essentially unchanged, the relaxed scenario exposes alternative refined optima that the strict scenario discards by construction. The C_1 – C_{70+} feed on the Pellegrini domain is the clearest illustration. Candidate 2 of Table 3.2 (T = 360 °C, P = 35 bar, H₂/wax = 0.150, WHSV = 3 kg_{n-C.in} kg_{cat}⁻¹ h⁻¹) reaches 51.21 wt-% of C_8 – C_{16} at the lower bound of the pressure range, against 52.03 wt-% at P = 60 bar for the strict-scenario optimum (candidate 3 of the same table). The two points trade off a 0.8 wt-% yield penalty against a 25 bar pressure reduction at the same WHSV and the same H₂/wax ratio. This is consistent with the kinetic structure of the model exposed in Section 3.2.2.3.1: lowering the pressure decreases p_{H_2} in the ADSORB denominator and accelerates the reaction rates. Operationally, candidate 2 is markedly less demanding: the lower pressure reduces compression duty for the hydrogen co-feed, allows a thinner vessel wall and lowers the capital expenditure on both the reactor and the upstream compression train; the lower temperature also reduces the heating

duty required to bring the feed to reactor conditions. Whether the 0.8 wt-% yield concession is justified by these savings is a question for a techno-economic assessment, which lies beyond the scope of this work. The point retained here is that the strict scenario, by accepting only yields within 0.01 wt-% of the maximum, suppresses this trade-off entirely.

The third observation concerns the systematic rejection of the cheapest candidate by the thermodynamic check. Across the two scenarios and the four feed/domain combinations, six of the eight most economical refined optima are rejected on VFRAC grounds and replaced by the next-ranked candidate. The pattern is identical in both scenarios and consistent with the kinetic model's insensitivity to H_2/wax established in Section 3.2.2.3.1: the search has no kinetic incentive to constrain the hydrogen ratio above its lower bound, so the economic ranking systematically promotes the lowest- H_2/wax candidate, which is also the one most likely to violate the vapour-fraction floor. The thermodynamic check is therefore not a marginal verification but the dominant filter that fixes the hydrogen ratio of every recommended optimum.

3.3.3 Synthesis

Table 3.3 collects the recommended operating conditions for the eight cases and forms the basis of the synthesis that follows.

Table 3.3: Recommended operating conditions for the eight cases explored.

Scenario	Feed	Domain	T (°C)	P (bar)	WHSV ($\text{kg}_{\text{n-C}_{16}}/\text{kg}_{\text{cat}}\text{h}^{-1}$)	H_2/wax	$C_8\text{--}C_{16}$ (wt-%)	$X_{C_{17+}}$ (%)	VFRAC	Rej.
Strict	$C_1\text{--}C_{70+}$	Pellegrini	369.3	60.0	3.00	0.150	52.03	55.9	0.984	0
Strict	$C_1\text{--}C_{70+}$	Extended	375.0	80.0	3.23	0.103	52.54	56.2	0.963	1
Strict	C_{17+}	Pellegrini	375.0	50.1	3.00	0.115	38.86	67.5	0.967	1
Strict	C_{17+}	Extended	375.0	43.0	3.45	0.139	38.86	67.5	0.978	0
Relaxed	$C_1\text{--}C_{70+}$	Pellegrini	360.0	35.0	3.00	0.150	51.21	54.5	0.990	1
Relaxed	$C_1\text{--}C_{70+}$	Extended	375.0	80.0	3.15	0.142	52.54	56.2	0.978	2
Relaxed	C_{17+}	Pellegrini	375.0	60.0	2.48	0.133	38.86	67.5	0.969	1
Relaxed	C_{17+}	Extended	375.0	37.0	4.00	0.150	38.86	67.5	0.984	2

Five observations can be retained from this study.

First, the C_{17+} feed makes substantially better use of the reactor than the full feed. The full $C_1\text{--}C_{70+}$ feed delivers a higher absolute outlet $C_8\text{--}C_{16}$ fraction, but this comparison is misleading because the full feed inherits 33.8 wt-% of $C_8\text{--}C_{16}$ from the upstream Fischer-Tropsch reactor. The net kerosene production of the hydrocracking step itself is therefore approximately 18 wt-% for the full feed against 39 wt-% for the C_{17+} feed, more than a twofold increase per unit mass of processed feedstock. The C_{17+} configuration also yields a higher heavy-end conversion ($X_{C_{17+}} \approx 67.5\%$ against $\sim 56\%$ for the full feed), since the absence of feed-borne kerosene removes the overcracking constraint that limits the operable temperature in the full-feed case. These conversion levels remain within the 20–100 % window reported for Co-LTFT wax hydrocracking in the literature (Table 1.3) but lie in its lower-to-middle range. The practical price of the C_{17+} strategy is the introduction of an upstream separation unit that withdraws $C_1\text{--}C_{16}$ from the Fischer-Tropsch effluent before the hydrocracker, the cost and energy demand of which lie beyond the scope of the present work but should be evaluated before this configuration is preferred in practice. It should further be noted that this C_{17+} upstream separation is assumed ideal throughout this work. The two configurations therefore correspond to the two limiting cases of upstream separation: a perfectly selective separation, for which only the C_{17+} heavy ends are fed and cracked, and the absence of any separation, for which the feed-borne $C_8\text{--}C_{16}$ fraction is exposed to overcracking. Any physically attainable separation lies between these two limits.

Secondly, the extended domain offers a marginal yield gain at the price of a complete loss of calibration support. On every case where both domains are explored, the search converges to a refined optimum on the extended domain that improves the yield by less than 0.5 wt-% relative to the Pellegrini-domain result. The extended optimum systematically sits at $T = 375\text{ }^\circ\text{C}$ (upper bound of both domains) and at

the upper bound of either pressure or WHSV, indicating that the solver is exploiting the wider exploration window primarily to push the temperature-pressure-WHSV combination further into the active regime. The corresponding gain is real within the model, but it is obtained at operating conditions outside the range where the kinetic parameters were validated, and there is no empirical basis on which to assess the reliability of the prediction in that region. For the purposes of process design, the Pellegrini-domain recommendation should therefore be preferred over the extended-domain recommendation whenever both produce comparable yields, since it carries an explicit kinetic basis.

Thirdly, boundary solutions dominate the recommendations. Seven of the eight recommended optima reported in Table 3.3 sit on at least two of the four domain boundaries simultaneously. The temperature is at the upper bound in six of the eight cases; the WHSV sits at its domain upper bound in four cases (at $3 \text{ kg}_{\text{n-C},\text{in}} \text{ kg}_{\text{cat}}^{-1} \text{ h}^{-1}$ for the three Pellegrini-domain optima and at $4 \text{ kg}_{\text{n-C},\text{in}} \text{ kg}_{\text{cat}}^{-1} \text{ h}^{-1}$ for one extended-domain optimum); and the hydrogen ratio is at or close to 0.15 kg/kg in six cases. These boundary locations must be interpreted with care. A corner solution does not establish that the true process optimum lies exactly at the edge of the explored window; it establishes only that, within the kinetic model taken in isolation, nothing penalises moving further in that direction. The domain bounds defined in Section 2.4.1 are not themselves part of the kinetic model: they encode constraints that the model does not represent. The temperature ceiling reflects the thermal stability of the catalyst and the range over which the kinetic parameters were calibrated; the pressure ceiling reflects mechanical and compression limits; and the WHSV and hydrogen-ratio bounds reflect the operating window documented in the literature for Co-LTFT wax hydrocracking. The kinetics alone would push the operating point beyond these edges, but crossing them would expose the reactor to phenomena that lie outside the scope of the present description, such as catalyst deactivation, mechanical limitations, or operation in a regime where the calibrated parameters are no longer reliable. The recommended optima must therefore be read as the best points within a deliberately bounded window, not as unconstrained process optima; releasing any bound would require a model that explicitly accounts for the constraint that bound represents.

Fourthly, the strict and relaxed scenarios differ less than expected on the recommendation, more than expected on the trade-off space. On the recommended optimum itself, the two scenarios deliver virtually identical operating conditions in the four cases involving the C_{17+} feed and in the extended domain for the full feed, because the dominant optimum lies well inside the relaxed iso-yield set in those cases. The only case where the relaxed scenario produces a substantially different recommended optimum is the $\text{C}_1\text{--}\text{C}_{70+}$ feed on the Pellegrini domain, where candidate #2 of Table 3.2 accesses a low-pressure, lower-temperature operating point at the cost of $0.8 \text{ wt-\%}$ of yield. The relaxed scenario therefore has no material impact on the recommended optimum in most cases, but it exposes alternative equivalent operating regimes that would have been suppressed by the strict tolerance, and which deserve to be retained in any subsequent techno-economic analysis where capital expenditure trade-offs can be quantified explicitly.

Fifthly, the thermodynamic check is the dominant filter on the hydrogen ratio. Combining the strict and relaxed scenarios across the eight feed/domain combinations, the procedure rejected post-hoc eight candidates on VFRAC grounds (column *Rej.* of Table 3.3) before reaching the recommended optimum. In every case, the rejected candidates featured a hydrogen ratio at the lower bound 0.06 kg/kg , and the recommendation eventually settled on a substantially higher ratio (0.103 to 0.150 kg/kg) selected purely to keep the inlet vapour fraction above the 96% admissibility floor. This pattern is not an artefact of the procedure: it reflects the structural insensitivity of the partial-pressure kinetic model to H_2/wax established in Section 3.2.2.3.1, which leaves the ratio free as far as the yield is concerned. The reported hydrogen ratios should consequently be understood as thermodynamically admissible defaults rather than as kinetically optimal values; a fugacity-based model in which H_2/wax controls the vapour-liquid distribution of the heavy paraffins would likely return a stronger kinetic incentive to operate at higher hydrogen ratios.

3.3.4 Sensitivity of the optimum to the feed composition

The kinetic parameters of the model were fitted to the Pellegrini et al. (2007) feedstock and were not refitted to the ULiège stream; the optima of Sections 3.3.1 and 3.3.2 are therefore obtained by applying

the model outside the composition on which it was calibrated, and must be read with this caveat. A direct way to evaluate the effect of this extrapolation is to repeat the optimum search on the Pellegrini feed itself, for which the model is used strictly in-domain, and to compare the result with the ULiège optimum. The comparison is carried out in the strict-precision scenario over the Pellegrini calibration domain, the setting in which the in-domain reference is most meaningful. The results are reported in Table 3.4.

Table 3.4: Recommended operating conditions for the ULiège and Pellegrini feeds, strict-precision scenario over the Pellegrini calibration domain.

Feed	No.	T (°C)	P (bar)	WHSV	H ₂ /wax	C ₈ –C ₁₆ (wt-%)	X _{C₁₇+} (%)	VFRAC	Validity
ULiège C ₁ –C ₇₀ +	1	369.3	60.0	3.00	0.150	52.03	55.9	0.984	Valid
Pellegrini	1	375.0	40.5	3.00	0.088	51.63	59.0	0.970	Valid (recommended)
	2	375.0	41.6	2.80	0.150	51.63	59.0	0.986	Valid
	3	375.0	60.0	1.94	0.150	51.63	59.0	0.981	Valid
	4	375.0	40.3	2.89	0.150	51.63	59.0	0.986	Valid
	5	375.0	51.3	2.27	0.150	51.63	59.0	0.984	Valid

The maximum attainable C₈–C₁₆ yield is almost identical for the two feeds: 52.03 wt-% for the ULiège stream against 51.63 wt-% for the Pellegrini feed, a difference of only 0.4 wt-%. The kerosene cut predicted by the model is thus essentially insensitive to this change of feed, both streams being closely comparable FT waxes. Within that scope, the ULiège value of 52 wt-% can be read as a robust order-of-magnitude estimate rather than an artefact of the extrapolation.

The location of the optimum, by contrast, must be compared with more care, because it is not a single sharp point. For the Pellegrini feed the procedure returns not one optimum but a plateau of five operating points that all reach 51.63 wt-% within the numerical precision of the integration; among them the economic criterion selects the point with the lowest hydrogen ratio (H₂/wax = 0.088). Along this plateau the temperature stays pinned at the upper bound (375 °C), while pressure, WHSV and the hydrogen ratio trade off at essentially constant yield. The single ULiège optimum (T = 369.3 °C, P = 60 bar, WHSV = 3.00, H₂/wax = 0.150) lies within this same high-yield region. The large apparent gap between the two recommended points (P = 60 vs 40.5 bar, H₂/wax = 0.150 vs 0.088) does not, however, correspond to a real difference in the optimal conditions: because the yield is flat over pressure and the hydrogen ratio, the economic tie-breaker selects widely separated coordinates on the same plateau, so the gap reflects the non-uniqueness of the selected point rather than a displacement of the optimum.

The only coordinate that shifts in a way not explained by the flatness of the response is the optimal temperature, and this shift has a clear physical origin. The ULiège stream already carries 33.8 wt-% of C₈–C₁₆ inherited from the upstream Fischer–Tropsch reactor, so an interior temperature (369.3 °C) is preferred: raising it further would crack this feed-borne kerosene faster than fresh kerosene is produced from the heavy ends (Section 3.2.2.1). As shown in Figure 3.17, the Pellegrini feed is heavier and contains less pre-formed kerosene (29.3 wt-%); with less kerosene to over-crack, the model drives the temperature to its upper bound (375 °C) and converts a larger share of the heavy ends (X_{C₁₇+} = 59.0% against 55.9% for the ULiège feed). The two feeds thus reach a comparable kerosene yield by different routes: the ULiège feed by preserving inherited kerosene at a moderated temperature, the Pellegrini feed by cracking its heavier ends more extensively. This is consistent with the larger amount of convertible heavy material in the Pellegrini feed. The thermodynamic admissibility criterion is satisfied in both cases (VFRAC ≥ 0.96).

In summary, both the achievable yield and the qualitative shape of the response (a yield governed mainly by temperature and only weakly sensitive to pressure, WHSV and the hydrogen ratio) are preserved across the two feeds; only the optimal temperature shifts, in a direction that the difference in feed composition explains. The weak sensitivity of the predicted optimum to this change of feed is an encouraging indication that the trends identified for the ULiège stream are not artefacts of applying the model outside its calibration range. Definitive confirmation would nonetheless require experimental hydrocracking data for the ULiège feed and a refit of the kinetic parameters to it.

Chapter 4

Conclusions and perspectives

4.1 Conclusions

Aviation is among the hardest sectors to decarbonise: long-haul flight relies on the high energy density of liquid hydrocarbons, for which neither battery-electric nor hydrogen propulsion offers a credible substitute on the relevant timescale. Producing those hydrocarbons from non-fossil carbon rather than from crude oil is therefore the central challenge, and the Fischer-Tropsch Power-to-Liquids route addresses it directly by converting captured CO₂ and renewable electricity into drop-in synthetic kerosene through a chain of catalytic conversion steps. It is within this framework that the University of Liège has been developing, through the Neutral-Kero-Lime project, an integrated Power-to-Kerosene chain in which the upstream reverse water-gas shift and cobalt-based low-temperature Fischer-Tropsch units had already been modelled by A. Rouxhet and A. Morales [1, 2]. The Fischer-Tropsch effluent, however, is weighted predominantly towards the heavy C₁₇₊ wax fraction, and no model of the hydrocracking unit required to upgrade this fraction into the C₈-C₁₆ kerosene cut had been integrated into the chain. The objective of the present work was therefore to model and simulate this hydrocracking step in Python, to identify the operating conditions that maximise the C₈-C₁₆ mass fraction at the reactor outlet, and to derive those conditions into an indicative reactor geometry.

Reproducing the experimental hydrocracking behaviour required a kinetic description. Both a stoichiometric reactor and a thermodynamic-equilibrium reactor, the two simplified descriptions most commonly retained when a hydrocracking unit is integrated into a broader process flowsheet, were shown to be structurally unable to reproduce the experimental outlet distribution: hydrocracking under the conditions investigated is kinetically rather than thermodynamically controlled. The all-components LHHW model of Pellegrini et al. (2007) [3] was retained on this basis, because it resolves the product stream at the component level, a resolution required to track the C₈-C₁₆ cut without committing to a fixed lump definition. Reconstructing it in Python required overcoming two omissions of the original publication. The inlet distribution, absent from the article, was identified by hypothesis testing as the component-level C₄-C₇₀ distribution of Pellegrini et al. (2008) [4], which reproduced the published lumped-model curves with $R^2 \approx 0.998$ -0.999. The unpublished final kinetic parameters were then recovered through a Tikhonov-regularised estimation procedure, reducing the sum of squared errors against the published model curves from 1.58×10^{-3} to 1.44×10^{-5} while remaining physically anchored to the reference functional forms. The reproduction is anchored to Pellegrini's published model rather than to an independent experimental dataset, and the residual underestimation of iso-paraffins in the kerosene range is a structural consequence of the midpoint cracking stoichiometry that no parameter adjustment can remove, as the authors themselves had already identified.

The validated model was then applied to the ULiège Fischer-Tropsch hydrocarbon stream after the olefin-hydrogenation pre-treatment. The structured four-phase search identified, for the full C₁-C₇₀₊ feed over the Pellegrini calibration domain, a recommended operating point at $T = 369.3^\circ\text{C}$, $P = 60\text{ bar}$, $\text{WHSV} = 3.00\text{ kg}_{\text{n-C},\text{in}}\text{ kg}_{\text{cat}}^{-1}\text{ h}^{-1}$ and $\text{H}_2/\text{wax} = 0.150\text{ kg/kg}$, yielding 52.0 wt-% of C₈-C₁₆ at the reactor

outlet for a C_{17+} mass conversion of 55.9 %. Three of the four coordinates lie on the upper boundary of the calibration domain, so this point is a corner solution rather than an interior optimum, and the model response would continue to climb beyond the calibrated range. Of the 52 wt-% of C_8 - C_{16} measured at the outlet, only +18 wt-% is genuinely produced by the hydrocracking step itself, the remainder being inherited unchanged from the feed; the broader analysis carried out over the eight cases combining two feed streams, two exploration domains and two precision scenarios further established that processing only the C_{17+} fraction nearly doubles the net kerosene production per unit mass of feed (39 wt-% against 18 wt-%), at the cost of an additional upstream separation unit. Extending the exploration beyond the calibration domain improves the yield only marginally (< 0.5 wt-%) while forfeiting all empirical support for the prediction. Finally, the thermodynamic admissibility criterion (inlet molar vapour fraction ≥ 96 %), and not the kinetics, is the dominant filter fixing the hydrogen-to-wax ratio, a direct consequence of the partial-pressure assumption, under which the model is structurally insensitive to that variable.

Translated into a reactor geometry, the recommended operating point corresponds, at the pilot-scale throughput of 0.114 kg/h, to an indicative cylindrical packed bed of approximately 26.5 mm internal diameter and 132 mm bed height, holding close to 38 g of bifunctional catalyst (0.6 wt-% Pt on amorphous silica-alumina). This geometry satisfies the plug-flow hydrodynamic criteria but remains an order-of-magnitude estimate rather than a construction-grade design.

Several caveats accompany these results. The kinetic parameters were calibrated on the Pellegrini C_4 - C_{70} feedstock and not refitted to the ULiège stream, so the reported yields and operating conditions remain order-of-magnitude indications. To evaluate the effect of this extrapolation, the optimum search was also run on the Pellegrini calibration feed itself, for which the model operates strictly in-domain. The achievable kerosene yield (51.6 against 52.0 wt-% for the ULiège feed) and the weak sensitivity of that yield to pressure, WHSV and the hydrogen ratio near the optimum were both preserved; the only appreciable difference lay in the optimal temperature, which sits at the upper bound for the heavier Pellegrini feed but at an interior value (369.3 °C) for the ULiège stream, whose higher content of feed-borne kerosene a moderated temperature protects from over-cracking. This consistency is reassuring as to the transferability of the trends from one feed to another, though it does not amount to a validation on the ULiège stream. Two structural limitations of the underlying model further bound their interpretation: the midpoint cracking stoichiometry, which underestimates the iso-paraffin yield in the kerosene range and produces a spurious accumulation at C_5 , and the partial-pressure assumption, which renders the model insensitive to the H_2 /wax ratio. Both have been addressed in later contributions of the same research series (the breakage probability function of Gamba et al. [6, 7] and the vapour-liquid equilibrium framework of Pellegrini et al. (2008) [4]), and constitute the natural starting point of the perspectives outlined in the next section.

Overall, the work delivers a reproducible, component-level kinetic model of the LTFT wax hydrocracker and a transferable methodology for locating its kerosene-maximising operating conditions, thereby providing the previously missing block of the ULiège Power-to-Kerosene chain.

4.2 Perspectives

Several directions follow naturally from this work, most directly to address the structural limitations of the model identified in Chapter 3.

The first concerns the kinetic model itself. The Pellegrini-Gambaro series demonstrated that agreement with experimental data improves markedly once two simplifying hypotheses retained in this work are relaxed: the partial-pressure assumption, replaced by a vapour-liquid equilibrium treatment in which fugacities enter the rate expressions [4], and the midpoint cracking stoichiometry, replaced by the breakage probability function [6, 7]. Implementing both would remove the insensitivity of the model to the H_2 /wax ratio and the underestimation of iso-paraffins in the kerosene range, together with the associated spurious accumulation at C_5 . The fugacity-based framework could not be stabilised in Python within the time available; revisiting this implementation, possibly with a more robust thermodynamic library, is therefore the most immediate

extension.

Beyond the reactor itself, several surrounding units remain to be modelled before the hydrocracking step can be assessed within the complete Power-to-Kerosene chain. An upstream separation unit should remove the fraction lighter than C_{17} from the Fischer-Tropsch effluent so that only the heavy C_{17+} fraction enters the hydrocracker; this feed-splitting configuration was shown in Section 3.3 to nearly double the net kerosene production of the step. The reactive stream must also be freed of the non-hydrocarbon species carried over from Fischer-Tropsch synthesis (water, CO_2 and CO), whose removal was assumed but not described in this work. The stream leaving the flash separator ($150\text{ }^{\circ}C$, 20 bar) must be brought to reactor operating conditions through a pump followed by a heat exchanger, the modelling of which is a prerequisite for any energy assessment of the process. On the product side, the reactor effluent leaves the C_{17+} feed only partially converted (55.9 wt-% at the recommended point), so a downstream flash separator recovering the unconverted C_{17+} wax and recycling it to the reactor inlet would raise the overall kerosene yield.

Two further extensions would consolidate the integration of this work into the ULiège process chain. Re-implementing the hydrocracking model in Aspen Custom Modeler would combine the equation-based flexibility demonstrated here with the thermodynamic database of Aspen Plus, enabling direct coupling with the upstream models of A. Rouxhet and A. Morales. Once the surrounding units are in place, a techno-economic study coupled with an energy-integration analysis would translate the reactor-scale results into a meaningful process-level assessment of the Power-to-Kerosene route.

The most impactful perspective, however, remains experimental. The kinetic parameters used in this work were calibrated on the C_4 - C_{70} feedstock of Pellegrini et al. and could not be refitted to the ULiège syncrude in the absence of any dedicated experimental dataset; the present results are therefore those of a published model applied by extrapolation, not of a model validated against the feed it is meant to describe. Acquiring hydrocracking measurements on a feed representative of the ULiège Fischer-Tropsch stream would allow the kinetic parameters to be re-estimated and the model to be genuinely validated, thereby converting the order-of-magnitude indications reported in this work into quantitative design values.

Appendices

Table A.1: Typical syncrude compositions for Fe-HTFT, Fe-LTFT and Co-LTFT Fischer-Tropsch syntheses [11].

Product fraction	Carbon range	Compound class	Syncrude composition ^a (mass%)		
			Fe-HTFT	Fe-LTFT	Co-LTFT
Tail gas	C ₁	Alkane	12.7	4.3	5.6
	C ₂	Alkene	5.6	1.0	0.1
		Alkane	4.5	1.0	1.0
LPG	C ₃ -C ₄	Alkene	21.2	6.0	3.4
		Alkane	3.0	1.8	1.8
Naphtha	C ₅ -C ₁₀	Alkene	25.8	7.7	7.8
		Alkane	4.3	3.3	12.0
		Aromatic	1.7	0	0
		Oxygenate	1.6	1.3	0.2
Distillate	C ₁₁ -C ₂₂	Alkene	4.8	5.7	1.1
		Alkane	0.9	13.5	20.8
		Aromatic	0.8	0	0
		Oxygenate	0.5	0.3	0
Residue/wax	C ₂₂ +	Alkene	1.6	0.7	0
		Alkane	0.4	49.2	44.6
		Aromatic	0.7	0	0
		Oxygenate	0.2	0	0
Aqueous product	C ₁ -C ₅	Alcohol	4.5	3.9	1.4
		Carbonyl	3.9	0	0
		Carboxylic acid	1.3	0.3	0.2

^a Syncrude composition is expressed as the total mass of product from Fischer-Tropsch synthesis, excluding inert gases and water-gas-shift products. Zero indicates low concentration and not necessarily total absence.

Table A.2: Thermodynamic state (A.2a) and molar composition of the stream (A.2b) at the flash separator outlet, as provided by A. Rouxhet.

(a) Thermodynamic properties.

Property	Unit	Value
Phase	—	Liquid
Temperature	°C	150.0
Pressure	bar	20
Molar enthalpy	kJ kmol^{-1}	-308 130.55
Mass enthalpy	kJ kg^{-1}	-2 493.93
Enthalpy flow	kW	-0.09
Mass density	kg m^{-3}	665.01
Average molecular weight	g mol^{-1}	123.55
Total volume flow rate	L h^{-1}	0.1856
Total molar flow	kmol h^{-1}	1.0×10^{-3}

(b) Molar composition (mol%).

Non-hydrocarbon species					
CO	0.283	H ₂ O	50.171	CO ₂	0.888
H ₂	0.353				
n-Paraffins					
C ₁ H ₄	1.063	C ₉ H ₂₀	1.043	C ₁₇ H ₃₆	2.921
C ₂ H ₆	0.132	C ₁₀ H ₂₂	1.485	C ₁₈ H ₃₈	2.623
C ₃ H ₈	0.099	C ₁₁ H ₂₄	1.933	C ₁₉ H ₄₀	2.314
C ₄ H ₁₀	0.119	C ₁₂ H ₂₆	2.483	C ₂₀ H ₄₂	2.019
C ₅ H ₁₂	1.037	C ₁₃ H ₂₈	3.000	C ₂₁ H ₄₄	1.747
C ₆ H ₁₄	0.790	C ₁₄ H ₃₀	3.261	C ₂₂ H ₄₆	1.506
C ₇ H ₁₆	0.686	C ₁₅ H ₃₂	3.327	C ₂₃ H ₄₈	1.294
C ₈ H ₁₈	0.787	C ₁₆ H ₃₄	3.164	C ₂₄ H ₅₀	1.109
C ₂₅ H ₅₂	0.951	C ₂₇ H ₅₆	0.696	C ₂₉ H ₆₀	0.509
C ₂₆ H ₅₄	0.814	C ₂₈ H ₅₈	0.596	C ₃₀₊	3.013
Olefins					
C ₂ H ₄	0.012	C ₅ H ₁₀	0.562	C ₈ H ₁₆	0.103
C ₃ H ₆	0.157	C ₆ H ₁₂	0.288	C ₉ H ₁₈	0.111
C ₄ H ₈	0.098	C ₇ H ₁₄	0.148	C ₁₀₊	0.305

Table A.3: Kinetic parameter values used for the first model in Pellegrini et al. (2007) [3].

(a) Langmuir and equilibrium constants.

Lump	K_{Liso} (1/Pa)	K_{Ln} (1/Pa)	K_{eq}
I (C ₃)	—	5.46E-02	—
II (C ₇)	5.46E+00	5.47E+00	3.59E+03
III (C ₁₂)	5.48E+01	5.48E+01	3.78E+03
IV (C ₁₈)	9.46E+02	9.05E+02	3.86E+03
V (C ₄₇)	9.53E+07	7.19E+07	5.93E+03

(b) Frequency factors.

Lump	k_{isom} (kmol/h/kg)	k_{cr} (kmol/h/kg)
I (C ₃)	—	—
II (C ₇)	1.14E+23	4.48E+22
III (C ₁₂)	7.60E+24	4.50E+24
IV (C ₁₈)	9.47E+27	3.05E+26
V (C ₄₇)	2.80E+29	2.99E+27

(c) Activation energies.

Lump	E_{isom} (kJ/kmol)	E_{cr} (kJ/kmol)
I (C ₃)	—	—
II (C ₇)	1.61E+05	1.34E+05
III (C ₁₂)	1.68E+05	1.65E+05
IV (C ₁₈)	1.92E+05	1.85E+05
V (C ₄₇)	2.01E+05	1.93E+05

Table A.4: Molar fractions in the FT hydrocarbon stream after removal of C_1 - C_{16} .

Paraffin	x_i	Paraffin	x_i	Paraffin	x_i	Paraffin	x_i
C_1H_4	0	$C_{19}H_{40}$	1.05×10^{-1}	$C_{37}H_{76}$	6.60×10^{-3}	$C_{55}H_{112}$	3.98×10^{-4}
C_2H_6	0	$C_{20}H_{42}$	9.13×10^{-2}	$C_{38}H_{78}$	5.65×10^{-3}	$C_{56}H_{114}$	3.41×10^{-4}
C_3H_8	0	$C_{21}H_{44}$	7.90×10^{-2}	$C_{39}H_{80}$	4.83×10^{-3}	$C_{57}H_{116}$	2.92×10^{-4}
C_4H_{10}	0	$C_{22}H_{46}$	6.81×10^{-2}	$C_{40}H_{82}$	4.13×10^{-3}	$C_{58}H_{118}$	2.49×10^{-4}
C_5H_{12}	0	$C_{23}H_{48}$	5.85×10^{-2}	$C_{41}H_{84}$	3.54×10^{-3}	$C_{59}H_{120}$	2.13×10^{-4}
C_6H_{14}	0	$C_{24}H_{50}$	5.02×10^{-2}	$C_{42}H_{86}$	3.03×10^{-3}	$C_{60}H_{122}$	1.83×10^{-4}
C_7H_{16}	0	$C_{25}H_{52}$	4.30×10^{-2}	$C_{43}H_{88}$	2.59×10^{-3}	$C_{61}H_{124}$	1.56×10^{-4}
C_8H_{18}	0	$C_{26}H_{54}$	3.68×10^{-2}	$C_{44}H_{90}$	2.22×10^{-3}	$C_{62}H_{126}$	1.34×10^{-4}
C_9H_{20}	0	$C_{27}H_{56}$	3.15×10^{-2}	$C_{45}H_{92}$	1.90×10^{-3}	$C_{63}H_{128}$	1.14×10^{-4}
$C_{10}H_{22}$	0	$C_{28}H_{58}$	2.69×10^{-2}	$C_{46}H_{94}$	1.62×10^{-3}	$C_{64}H_{130}$	9.78×10^{-5}
$C_{11}H_{24}$	0	$C_{29}H_{60}$	2.30×10^{-2}	$C_{47}H_{96}$	1.39×10^{-3}	$C_{65}H_{132}$	8.37×10^{-5}
$C_{12}H_{26}$	0	$C_{30}H_{62}$	1.97×10^{-2}	$C_{48}H_{98}$	1.19×10^{-3}	$C_{66}H_{134}$	7.16×10^{-5}
$C_{13}H_{28}$	0	$C_{31}H_{64}$	1.68×10^{-2}	$C_{49}H_{100}$	1.02×10^{-3}	$C_{67}H_{136}$	6.13×10^{-5}
$C_{14}H_{30}$	0	$C_{32}H_{66}$	1.44×10^{-2}	$C_{50}H_{102}$	8.69×10^{-4}	$C_{68}H_{138}$	5.24×10^{-5}
$C_{15}H_{32}$	0	$C_{33}H_{68}$	1.23×10^{-2}	$C_{51}H_{104}$	7.43×10^{-4}	$C_{69}H_{140}$	4.48×10^{-5}
$C_{16}H_{34}$	0	$C_{34}H_{70}$	1.05×10^{-2}	$C_{52}H_{106}$	6.36×10^{-4}	$C_{70}H_{142}$	2.66×10^{-4}
$C_{17}H_{36}$	1.32×10^{-1}	$C_{35}H_{72}$	9.02×10^{-3}	$C_{53}H_{108}$	5.44×10^{-4}		
$C_{18}H_{38}$	1.19×10^{-1}	$C_{36}H_{74}$	7.72×10^{-3}	$C_{54}H_{110}$	4.66×10^{-4}	Total	1.00

Table A.5: Initial molar-fraction distribution x_i of n-paraffins C_nH_{2n+2} from Pellegrini et al. (2008) [4].

Paraffin	x_i	Paraffin	x_i	Paraffin	x_i	Paraffin	x_i
C_1H_4	0	$C_{19}H_{40}$	3.70×10^{-2}	$C_{37}H_{76}$	6.43×10^{-3}	$C_{55}H_{112}$	1.11×10^{-3}
C_2H_6	0	$C_{20}H_{42}$	3.43×10^{-2}	$C_{38}H_{78}$	5.91×10^{-3}	$C_{56}H_{114}$	1.13×10^{-3}
C_3H_8	0	$C_{21}H_{44}$	3.18×10^{-2}	$C_{39}H_{80}$	5.50×10^{-3}	$C_{57}H_{116}$	9.31×10^{-4}
C_4H_{10}	3.83×10^{-4}	$C_{22}H_{46}$	2.97×10^{-2}	$C_{40}H_{82}$	5.08×10^{-3}	$C_{58}H_{118}$	8.43×10^{-4}
C_5H_{12}	1.03×10^{-3}	$C_{23}H_{48}$	2.81×10^{-2}	$C_{41}H_{84}$	4.62×10^{-3}	$C_{59}H_{120}$	8.04×10^{-4}
C_6H_{14}	3.27×10^{-3}	$C_{24}H_{50}$	2.67×10^{-2}	$C_{42}H_{86}$	4.15×10^{-3}	$C_{60}H_{122}$	7.65×10^{-4}
C_7H_{16}	1.97×10^{-2}	$C_{25}H_{52}$	2.55×10^{-2}	$C_{43}H_{88}$	3.78×10^{-3}	$C_{61}H_{124}$	4.75×10^{-4}
C_8H_{18}	3.20×10^{-2}	$C_{26}H_{54}$	2.30×10^{-2}	$C_{44}H_{90}$	3.45×10^{-3}	$C_{62}H_{126}$	4.84×10^{-4}
C_9H_{20}	4.44×10^{-2}	$C_{27}H_{56}$	2.07×10^{-2}	$C_{45}H_{92}$	3.12×10^{-3}	$C_{63}H_{128}$	5.39×10^{-4}
$C_{10}H_{22}$	5.45×10^{-2}	$C_{28}H_{58}$	1.83×10^{-2}	$C_{46}H_{94}$	2.83×10^{-3}	$C_{64}H_{130}$	5.43×10^{-4}
$C_{11}H_{24}$	5.94×10^{-2}	$C_{29}H_{60}$	1.59×10^{-2}	$C_{47}H_{96}$	2.55×10^{-3}	$C_{65}H_{132}$	5.43×10^{-4}
$C_{12}H_{26}$	6.04×10^{-2}	$C_{30}H_{62}$	1.35×10^{-2}	$C_{48}H_{98}$	2.27×10^{-3}	$C_{66}H_{134}$	5.43×10^{-4}
$C_{13}H_{28}$	5.89×10^{-2}	$C_{31}H_{64}$	1.19×10^{-2}	$C_{49}H_{100}$	1.98×10^{-3}	$C_{67}H_{136}$	1.75×10^{-4}
$C_{14}H_{30}$	5.56×10^{-2}	$C_{32}H_{66}$	1.05×10^{-2}	$C_{50}H_{102}$	1.67×10^{-3}	$C_{68}H_{138}$	2.67×10^{-5}
$C_{15}H_{32}$	5.18×10^{-2}	$C_{33}H_{68}$	9.42×10^{-3}	$C_{51}H_{104}$	1.47×10^{-3}	$C_{69}H_{140}$	0
$C_{16}H_{34}$	4.70×10^{-2}	$C_{34}H_{70}$	8.66×10^{-3}	$C_{52}H_{106}$	1.29×10^{-3}	$C_{70}H_{142}$	0
$C_{17}H_{36}$	4.35×10^{-2}	$C_{35}H_{72}$	7.96×10^{-3}	$C_{53}H_{108}$	1.13×10^{-3}		
$C_{18}H_{38}$	4.09×10^{-2}	$C_{36}H_{74}$	7.20×10^{-3}	$C_{54}H_{110}$	9.93×10^{-4}		

Table A.6: Optimised parameters used to reproduce the Pellegrini et al. (2007) model.

Parameter	Unit	Initial	Optimised	Rel. dev. %
KLn_fact	Pa^{-1}	1×10^{-1}	4.347×10^{-1}	+334.7%
KLn_exp	—	3.9×10^{-1}	3.898×10^{-1}	−0.1%
KLiso_fact	Pa^{-1}	1×10^{-1}	1.078×10^{-1}	+7.8%
KLiso_exp	—	3.8×10^{-1}	3.855×10^{-1}	+1.4%
Keq_carre	—	1.1268	1.1268	+0.0%
Keq_lin	—	3.1914	3.1914	−0.0%
Keq_cst	—	3589.1	3589.14	+0.0%
kisom_fact	$\text{kmol h}^{-1} \text{kg}^{-1}$	5.00×10^{16}	5.265×10^{16}	+5.3%
kisom_exp	—	7.9241	8.021	+1.2%
kcr0_fact	$\text{kmol h}^{-1} \text{kg}^{-1}$	2.00×10^{18}	1.981×10^{18}	−1.0%
kcr_exp	—	5.7677	5.851	+1.5%
Eiso_fact	kJ kmol^{-1}	22253	22226	−0.1%
Eiso_cst	kJ kmol^{-1}	118350	119599	+1.1%
Ecr_fact	kJ kmol^{-1}	29806	29440	−1.2%
Ecr_cst	kJ kmol^{-1}	86006	87185	+1.4%

Table A.7: Detailed molar fractions of Fischer-Tropsch products provided by A. Rouxhet.

n_c	Paraffins $\text{C}_n\text{H}_{2n+2}$	Olefins C_nH_{2n}	After hydrogenation	n_c	Paraffins $\text{C}_n\text{H}_{2n+2}$	Olefins C_nH_{2n}	After hydrogenation
1	2.20×10^{-2}	0	2.20×10^{-2}	36	3.53×10^{-3}	1.32×10^{-8}	3.53×10^{-3}
2	2.73×10^{-3}	2.47×10^{-4}	2.97×10^{-3}	37	3.02×10^{-3}	8.28×10^{-9}	3.02×10^{-3}
3	2.06×10^{-3}	3.24×10^{-3}	5.30×10^{-3}	38	2.59×10^{-3}	5.20×10^{-9}	2.59×10^{-3}
4	2.46×10^{-3}	2.03×10^{-3}	4.49×10^{-3}	39	2.21×10^{-3}	3.27×10^{-9}	2.21×10^{-3}
5	2.15×10^{-2}	1.16×10^{-2}	3.31×10^{-2}	40	1.89×10^{-3}	2.05×10^{-9}	1.89×10^{-3}
6	1.64×10^{-2}	5.96×10^{-3}	2.23×10^{-2}	41	1.62×10^{-3}	1.29×10^{-9}	1.62×10^{-3}
7	1.42×10^{-2}	3.07×10^{-3}	1.73×10^{-2}	42	1.39×10^{-3}	8.09×10^{-10}	1.39×10^{-3}
8	1.63×10^{-2}	2.14×10^{-3}	1.84×10^{-2}	43	1.19×10^{-3}	5.08×10^{-10}	1.19×10^{-3}
9	2.16×10^{-2}	2.29×10^{-3}	2.39×10^{-2}	44	1.02×10^{-3}	3.19×10^{-10}	1.02×10^{-3}
10	3.07×10^{-2}	2.35×10^{-3}	3.31×10^{-2}	45	8.68×10^{-4}	2.01×10^{-10}	8.68×10^{-4}
11	4.00×10^{-2}	1.48×10^{-3}	4.15×10^{-2}	46	7.43×10^{-4}	1.26×10^{-10}	7.43×10^{-4}
12	5.14×10^{-2}	9.27×10^{-4}	5.23×10^{-2}	47	6.35×10^{-4}	7.91×10^{-11}	6.35×10^{-4}
13	6.21×10^{-2}	5.82×10^{-4}	6.27×10^{-2}	48	5.44×10^{-4}	4.97×10^{-11}	5.44×10^{-4}
14	6.75×10^{-2}	3.66×10^{-4}	6.79×10^{-2}	49	4.65×10^{-4}	3.12×10^{-11}	4.65×10^{-4}
15	6.89×10^{-2}	2.30×10^{-4}	6.91×10^{-2}	50	3.98×10^{-4}	1.96×10^{-11}	3.98×10^{-4}
16	6.55×10^{-2}	1.44×10^{-4}	6.56×10^{-2}	51	3.40×10^{-4}	1.23×10^{-11}	3.40×10^{-4}
17	6.05×10^{-2}	9.06×10^{-5}	6.06×10^{-2}	52	2.91×10^{-4}	7.73×10^{-12}	2.91×10^{-4}
18	5.43×10^{-2}	5.69×10^{-5}	5.43×10^{-2}	53	2.49×10^{-4}	4.86×10^{-12}	2.49×10^{-4}
19	4.79×10^{-2}	3.58×10^{-5}	4.79×10^{-2}	54	2.13×10^{-4}	3.05×10^{-12}	2.13×10^{-4}
20	4.18×10^{-2}	2.25×10^{-5}	4.18×10^{-2}	55	1.82×10^{-4}	1.92×10^{-12}	1.82×10^{-4}
21	3.62×10^{-2}	1.41×10^{-5}	3.62×10^{-2}	56	1.56×10^{-4}	1.20×10^{-12}	1.56×10^{-4}
22	3.12×10^{-2}	8.86×10^{-6}	3.12×10^{-2}	57	1.34×10^{-4}	7.56×10^{-13}	1.34×10^{-4}
23	2.68×10^{-2}	5.56×10^{-6}	2.68×10^{-2}	58	1.14×10^{-4}	4.75×10^{-13}	1.14×10^{-4}
24	2.30×10^{-2}	3.50×10^{-6}	2.30×10^{-2}	59	9.77×10^{-5}	2.98×10^{-13}	9.77×10^{-5}
25	1.97×10^{-2}	2.20×10^{-6}	1.97×10^{-2}	60	8.36×10^{-5}	1.87×10^{-13}	8.36×10^{-5}
26	1.68×10^{-2}	1.38×10^{-6}	1.68×10^{-2}	61	7.15×10^{-5}	1.18×10^{-13}	7.15×10^{-5}
27	1.44×10^{-2}	8.66×10^{-7}	1.44×10^{-2}	62	6.12×10^{-5}	7.39×10^{-14}	6.12×10^{-5}
28	1.23×10^{-2}	5.44×10^{-7}	1.23×10^{-2}	63	5.24×10^{-5}	4.64×10^{-14}	5.24×10^{-5}
29	1.05×10^{-2}	3.42×10^{-7}	1.05×10^{-2}	64	4.48×10^{-5}	2.92×10^{-14}	4.48×10^{-5}
30	9.01×10^{-3}	2.15×10^{-7}	9.01×10^{-3}	65	3.83×10^{-5}	1.83×10^{-14}	3.83×10^{-5}
31	7.71×10^{-3}	1.35×10^{-7}	7.71×10^{-3}	66	3.28×10^{-5}	1.15×10^{-14}	3.28×10^{-5}
32	6.60×10^{-3}	8.47×10^{-8}	6.60×10^{-3}	67	2.81×10^{-5}	7.23×10^{-15}	2.81×10^{-5}
33	5.64×10^{-3}	5.32×10^{-8}	5.64×10^{-3}	68	2.40×10^{-5}	4.54×10^{-15}	2.40×10^{-5}
34	4.83×10^{-3}	3.34×10^{-8}	4.83×10^{-3}	69	2.05×10^{-5}	2.85×10^{-15}	2.05×10^{-5}
35	4.13×10^{-3}	2.10×10^{-8}	4.13×10^{-3}	70+	1.22×10^{-4}	4.82×10^{-15}	1.22×10^{-4}
Total					9.63×10^{-1}	3.69×10^{-2}	1.00

Disclaimer on AI

The use of artificial intelligence in scientific research remains a controversial topic, as it is often associated with rushed work or outputs of lower quality. However, when used thoughtfully and transparently, I believe that AI can contribute to producing higher-quality work within constrained timeframes. In the present study, the AI tool Perplexity was used to identify relevant sources based on selected keywords. The source articles retrieved through this process were, of course, carefully read and critically analysed before being incorporated into the work.

Claude AI was used to support the development of the code for the optimisation of kinetic parameters presented in Section 2.3.2, as well as the algorithm used to determine the optimal operating conditions discussed in Section 2.4. Given the complexity of these algorithms, their implementation would likely have required several weeks without such support; with the assistance of this tool, it was possible to complete this work within a few days. Nevertheless, all outputs generated by Claude, including code and results, were subsequently reviewed, adjusted, corrected, verified, and properly documented before being used.

Claude was also used as a writing aid, particularly to help formulate ideas and select appropriate terminology in English. Artificial intelligence should, in my opinion, therefore be regarded as a tool that can improve efficiency and save time, but not as a substitute for the user's own critical thinking, understanding, and scientific judgement. I remain fully responsible for the content of this work, including the validity of the results, the interpretation of the findings, and the accuracy of the manuscript.

Bibliography

- [1] Fabrice Orban, Grégoire Léonard et Yoshie Yonekura : Belgium's next century saf / e-fuel ecosystem. 28 October 2021.
- [2] Alejandro Morales Perez, Antoine Rouxhet et Grégoire Léonard : Fischer-tropsch kinetics comparison and sensitivity analysis for a pilot-scale jet fuel production reactor. 2022.
- [3] L. Pellegrini, S. Bonomi, S. Gamba, V. Calemme et D. Molinari : The "all components hydrocracking model". *Chemical Engineering Science*, 62(18-20):5013–5020, 2007.
- [4] Laura A. Pellegrini, Simone Gamba, Vincenzo Calemme et Susi Bonomi : Modelling of hydrocracking with vapour–liquid equilibrium. *Chemical Engineering Science*, 63(17):4285–4291, 2008.
- [5] L. Pellegrini, S. Locatelli, S. Rasella, S. Bonomi et V. Calemme : Modeling of fischer–tropsch products hydrocracking. *Chemical Engineering Science*, 59(22-23):4781–4787, novembre 2004.
- [6] Simone Gamba, Laura A. Pellegrini, Vincenzo Calemme et Chiara Gambaro : Introduction of a break-age probability function in the hydrocracking reactor model. *Industrial amp; Engineering Chemistry Research*, 48(12):5656–5665, avril 2009.
- [7] Chiara Gambaro, Vincenzo Calemme, Daniele Molinari et Joeri Denayer : Hydrocracking of fischer-tropsch waxes: Kinetic modeling via lhhw approach. *AIChE Journal*, 57(3):711–723, février 2011.
- [8] German Environment Agency : Power-to-liquids: A scalable and sustainable fuel supply perspective for aviation. Rapport technique, Umweltbundesamt, Dessau-Roßlau, Germany, janvier 2022.
- [9] Guangrui Liu, Beibei Yan et Guanyi Chen : Technical review on jet fuel production. *Renewable and Sustainable Energy Reviews*, 25:59–70, 2013.
- [10] Ali Elwalily, Emma Verkama, Franz Mantei, Adiya Kaliyeva, Andrew Pounder, Jörg Sauer et Florian Nestler : Sustainable aviation fuel production via the methanol pathway: a technical review. *Sustainable Energy amp; Fuels*, 9(19):5151–5180, 2025.
- [11] Arno de Klerk : *Fischer-Tropsch Refining*. Wiley, 2011.
- [12] A DEKLERK : Hydroprocessing peculiarities of fischer–tropsch syncrude. *Catalysis Today*, 130(2-4):439–445, janvier 2008.
- [13] Julius Scherzer et A.J. Gruia : Hydrocracking science and technology, août 1996.
- [14] J. G. Speight et Baki Ozum, éditeurs. *Petroleum Refining Processes*. CRC Press, 1 édition, 2001.
- [15] Matthew A. Koen : Hydrocracking of long chain n-paraffins under fischer-tropsch conditions. Msc thesis, University of Cape Town, Cape Town, South Africa, août 2014.
- [16] K. Tasaka, Y. Tanaka et M. Iwama : Hydrocracking process and process for producing hydrocarbon oil, 2012.

- [17] Olga Pleyer, Petr Straka, Dan Vrtiška, Jiří Hájek et Radek Černý : Hydrocracking of fischer-tropsch wax. *Paliva*, page 26–33, mars 2020.
- [18] Homepage V2 — corporate.ryanair.com. <https://corporate.ryanair.com/page/68/?market=ie>. [Accessed 19-05-2026].
- [19] International Energy Agency (IEA) : Co2 emissions in aviation in the net zero scenario, 2000–2030, 2025. Licence: CC BY 4.0.
- [20] Stef Proost : What role for aviation in climate policy? *Journal of Shipping and Trade*, 11(1), avril 2026.
- [21] European Commission : ReFuelEU aviation. https://transport.ec.europa.eu/transport-modes/air/environment/refueleu-aviation_en. [Accessed 25-04-2026].
- [22] Ahmed Rufai Dahiru, Ari Vuokila et Mika Huuhtanen : Recent development in power-to-x: Part i - a review on techno-economic analysis. *Journal of Energy Storage*, 56:105861, décembre 2022.
- [23] Muhammad R. Usman : Hydrogen storage methods: Review and current status. *Renewable and Sustainable Energy Reviews*, 167:112743, 2022.
- [24] Thierry Poinot Marc Fontecave, Sébastien Candel : L'hydrogène aujourd'hui et demain. https://www.academie-sciences.fr/pdf/rapport/rapport_hydrogene_2024.pdf, 2024. [Accessed 23-04-2026].
- [25] M. Hermesmann et T.E. Müller : Green, turquoise, blue, or grey? environmentally friendly hydrogen production in transforming energy systems. *Progress in Energy and Combustion Science*, 90:100996, 2022.
- [26] Merve Ozturk et Ibrahim Dincer : A comprehensive review on power-to-gas with hydrogen options for cleaner applications. *International Journal of Hydrogen Energy*, 46(62):31511–31522, 2021.
- [27] S.A. Grigoriev, V.N. Fateev, D.G. Bessarabov et P. Millet : Current status, research trends, and challenges in water electrolysis science and technology. *International Journal of Hydrogen Energy*, 45(49):26036–26058, 2020. Progress in Hydrogen Production and Utilization.
- [28] Vincent Dieterich, Alexander Buttler, Andreas Hanel, Hartmut Spliethoff et Sebastian Fendt : Power-to-liquid via synthesis of methanol, dme or fischer–tropsch-fuels: a review. *Energy amp; Environmental Science*, 13(10):3207–3252, 2020.
- [29] International Energy Agency : Global hydrogen review 2024. Rapport technique, IEA, Paris, 2024.
- [30] Karim Ghaib et Fatima-Zahrae Ben-Fares : Power-to-methane: A state-of-the-art review. *Renewable and Sustainable Energy Reviews*, 81:433–446, 2018.
- [31] Ekkehard Schwab, Andrian Milanov, Stephan Andreas Schunk, Axel Behrens et Nicole Schödel : Dry reforming and reverse water gas shift: Alternatives for syngas production? *Chemie Ingenieur Technik*, 87(4):347–353, 2015.
- [32] C. Bouchy, G. Hastoy, E. Guillon et J. A. Martens : Fischer-tropsch waxes upgrading via hydrocracking and selective hydroisomerization. *Oil amp; Gas Science and Technology - Revue de l'IFP*, 64(1): 91–112, janvier 2009.
- [33] Alejandro Ayala-Cortés, Christian Di Stasi, Daniel Torres, José Luis Pinilla et Isabel Suelves : Upgrading fischer-tropsch waxes to produce transport fuels by catalytic hydrocracking/isomerization: A review. *Renewable and Sustainable Energy Reviews*, 215:115633, 2025.

- [34] Antoine Rouxhet : Valoriser le CO_2 : solutions d'ingénierie circulaire. <https://orbi.uliege.be/bitstream/2268/342846/1/Universit%C3%A9%20des%20Champs%202026%20-%20Valorisation%20du%20CO2%20-%20Antoine%20Rouxhet.pdf>, 2026. [Accessed 23-04-2026], Presentation as part of a PhD thesis.
- [35] Stefan Bube, Steffen Voß, Nils Bullerdiek et Martin Kaltschmitt : *Power-to-Chemicals: Defossilization of the Chemical Industry via PtX*, page 941–974. Springer Nature Switzerland, décembre 2024.
- [36] Jussi Ikäheimo, Juha Kiviluoma, Robert Weiss et Hannele Holttinen : Power-to-ammonia in future north european 100 % renewable power and heat system. *International Journal of Hydrogen Energy*, 43(36):17295–17308, 2018.
- [37] Sudipta Chatterjee, Indranil Dutta et Kuo-Wei Huang : *Power to formic acid*, page 169–210. Elsevier, 2021.
- [38] International Renewable Energy Agency (IRENA) : Renewable power-to-heat. https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2019/Sep/IRENA_Power-to-heat_2019.pdf?la=en&hash=524C1, 2019. [Accessed 25-04-2026].
- [39] Harold Schobert : *Middle distillate fuels*, page 256–280. Cambridge Series in Chemical Engineering. Cambridge University Press, 2013.
- [40] Andreas Meurer et Jürgen Kern : Fischer–tropsch synthesis as the key for decentralized sustainable kerosene production. *Energies*, 14(7):1836, mars 2021.
- [41] International Energy Agency (IEA) : Oil market report. Rapport technique 11, International Energy Agency, sep 2025.
- [42] Paul Bardon et Olivier Massol : Decarbonizing aviation with sustainable aviation fuels: Myths and realities of the roadmaps to net zero by 2050. *Renewable and Sustainable Energy Reviews*, 211: 115279, avril 2025.
- [43] IATA : SAF Production Growth Rate is Slowing Down, Essential to Correct Course Ahead of e-SAF Mandates. <https://www.iata.org/en/pressroom/2025-releases/2025-12-09-04/>, December 2025. [Accessed 26-04-2026].
- [44] A. Anuar, V.K. Undavalli, B. Khandelwal et S. Blakey : Effect of fuels, aromatics and preparation methods on seal swell. *The Aeronautical Journal*, 125(1291):1542–1565, avril 2021.
- [45] V. Calemme, S. Corraera, C. Perego, P. Pollesel et L. Pellegrini : Hydroconversion of fischer–tropsch waxes: Assessment of the operating conditions effect by factorial design experiments. *Catalysis Today*, 106(1-4):282–287, octobre 2005.
- [46] S.T. Sie, M.M.G. Senden et H.M.H. Van Wechem : Conversion of natural gas to transportation fuels via the shell middle distillate synthesis process (smds). *Catalysis Today*, 8(3):371–394, janvier 1991.
- [47] Dieter Leckel et Maurice Liwanga-Ehumbu : Diesel-selective hydrocracking of an iron-based fischer-tropsch wax fraction (c15c45) using a moo3-modified noble metal catalyst. *Energy amp; Fuels*, 20(6):2330–2336, 2006.
- [48] L.P. Dancuart, R. de Haan et A. de Klerk : Chapter 6 - processing of primary fischer-tropsch products. In André Steynberg et Mark Dry, éditeurs : *Fischer-Tropsch Technology*, volume 152 de *Studies in Surface Science and Catalysis*, pages 482–532. Elsevier, 2004.
- [49] Arno de Klerk : Thermal cracking of fischer-tropsch waxes. *Industrial amp; Engineering Chemistry Research*, 46(17):5516–5521, 2007.

- [50] Arno de Klerk : Fischer–tropsch refining: technology selection to match molecules. *Green Chemistry*, 10(12):1249, 2008.
- [51] Arno de Klerk, Garima Chauhan, Cibele Halmenschlager, Felix Link, Natalia Montoya Sánchez, Brian Gartley, Hanan E. M. El-Sayed, Ranjit Sehdev et Rick Lehoux : Sustainable aviation fuel: Pathways to fully formulated synthetic jet fuel via fischer–tropsch synthesis. *Energy Science amp; Engineering*, 12(2):394–409, janvier 2023.
- [52] Pasquale Licastro, Celeste Scicchitano, Luca Mendicino, Emanuel Muraca et Domenico Macriapos; : Techno-economic and life cycle assessment of fischer-tropsch and methanol-to-jet processes to produce saf from captured carbon dioxide and green hydrogen. 2025.
- [53] Harold Schobert : *Composition, classification, and properties of petroleum*, page 174–191. Cambridge Series in Chemical Engineering. Cambridge University Press, 2013.
- [54] Dieter Leckel : H₂s effects in base-metal-catalyzed hydrocracking of fischertropsch wax. *Energy amp; Fuels*, 23(5):2370–2375, mai 2009.
- [55] Fabiano A.N. Fernandes et Ulisses M. Teles : Modeling and optimization of fischer–tropsch products hydrocracking. *Fuel Processing Technology*, 88(2):207–214, février 2007.
- [56] Paul B. Weisz : *Polyfunctional Heterogeneous Catalysis*, page 137–190. Elsevier, 1962.
- [57] H. L. Coonradt et W. E. Garwood : Mechanism of hydrocracking. reactions of paraffins and olefins. *Industrial amp; Engineering Chemistry Process Design and Development*, 3(1):38–45, janvier 1964.
- [58] Laura A. Pellegrini, Simone Gamba, Susi Bonomi et Vincenzo Calemme : Equilibrium constants for isomerization of n-paraffins. *Industrial amp; Engineering Chemistry Research*, 46(16):5446–5452, 2007.
- [59] Klaus Möller, Philip le Grange et Carlo Accolla : A two-phase reactor model for the hydrocracking of fischertropsch-derived wax. *Industrial amp; Engineering Chemistry Research*, 48(8):3791–3801, mars 2009.
- [60] JENS WEITKAMP : *The Influence of Chain Length in Hydrocracking and Hydroisomerization of n-Alkanes*, page 1–27. AMERICAN CHEMICAL SOCIETY, 1975.
- [61] Christian Marcilly : Réactivité et modes de transformation des grandes familles d'hydrocarbures. In *Catalyse acido-basique, application au raffinage et à la pétrochimie*. Technip, Paris, 2003.
- [62] Yoni Austen : Study of separation and refining processes of effluents from a fischer-tropsch reactor for the synthesis of e-kerosene from co₂ and low-carbon electricity, juin 2024.
- [63] R. C. Archibald, B. S. Greensfelder, George Holzman et D. H. Rowe : Catalytic hydrocracking of aliphatic hydrocarbons. *Industrial amp; Engineering Chemistry*, 52(9):745–750, 1960.
- [64] S. Ernst J. Weitkamp : *Guidelines for Mastering the Properties of Molecular Sieves: Relationship between the Physicochemical Properties of Zeolitic Systems and Their Low Dimensionality*. Springer US, 1990. p. 343.
- [65] G. F. Froment : Kinetics of the hydroisomerization and hydrocracking of paraffins on a platinum containing bifunctional γ -zeolite. *Catalysis Today*, 1(4):455–473, 1987.
- [66] M. A. Baltanas, K. K. Van Raemdonck, G. F. Froment et S. R. Mohedas : Fundamental kinetic modeling of hydroisomerization and hydrocracking on noble metal-loaded faujasites. 1. rate parameters for hydroisomerization. *Industrial & Engineering Chemistry Research*, 28(7):899–910, 1989.

- [67] E. Vynckier et G. F. Froment : Modeling of the kinetics of complex processes based upon elementary steps. In G. Astarita et S. I. Sandler, éditeurs : *Kinetic and Thermodynamic Lumping of Multicomponent Mixtures*, pages 131–161. Elsevier, Amsterdam, 1991.
- [68] G. G. Martens, G. B. Marin, J. A. Martens, P. A. Jacobs et G. V. Baron : A fundamental kinetic model for hydrocracking of C_8 to C_{12} alkanes on pt/us-y zeolites. *Journal of Catalysis*, 195(2):253–267, 2000.
- [69] Michael T. Klein et Gang Hou : *Mechanistic Kinetic Modeling of Heavy Paraffin Hydrocracking*, page 635–652. Springer New York.
- [70] J.-M. Schweitzer, P. Galtier et D. Schweich : A single-events kinetic model for the hydrocracking of paraffins in a three-phase reactor. *Chemical Engineering Science*, 54(13–14):2441–2452, 1999.
- [71] M. Steijns, G. F. Froment, P. Jacobs, J. Uytterhoeven et J. Weitkamp : Hydroisomerization and hydrocracking. 2. product distributions from n-decane and n-dodecane. 20(4):654–660, 1981.
- [72] M. Irfan Hosukoglu, Mustafa Karakaya et Ahmet K. Avci : Modeling and simulation of hydrocracking of fischer–tropsch hydrocarbons in a catalytic microchannel reactor. *Industrial and Engineering Chemistry Research*, 51(26):8913–8921, avril 2012.
- [73] Davide Selvatico, Andrea Lanzini et Massimo Santarelli : Low temperature fischer–tropsch fuels from syngas: Kinetic modeling and process simulation of different plant configurations. *Fuel*, 186:544–560, décembre 2016.
- [74] Gi Hoon Hong, Juwon Lee, Youngtak Cho et Sungwon Hwang : Techno-economic analysis of CO_2 /steam co-electrolysis process and synfuel production process coupled with steel manufacturing process. *Korean Journal of Chemical Engineering*, 40(4):740–753, mars 2023.
- [75] Solomon M. Jacob, Benjamin Gross, Sterling E. Voltz et Vern W. Weekman : A lumping and reaction scheme for catalytic cracking. *AIChE Journal*, 22(4):701–713, 1976.
- [76] Giorgio S. Soave : Estimation of the critical constants of heavy hydrocarbons for their treatment by the soave–redlich–kwong equation of state. *Fluid Phase Equilibria*, 143(1-2):29–39, janvier 1998.
- [77] S. Mohanty, D. Kunzru et D.N. Saraf : Hydrocracking: a review. *Fuel*, 69(12):1467–1473, décembre 1990.
- [78] Fredrik Nyholm, Sami Toppinen et Henrik Saxén : Sustainable power-to-liquids aviation fuels: Modelling and comparison of two fischer–tropsch upgrading process concepts. *Energy Conversion and Management*, 342:120153, 2025.
- [79] Bo Lai, Youliang Shi et Li Xu : Method of hydrotreatment of fischer–tropsch synthesis products, 2019. Application No. 15/673,439; continuation-in-part of PCT/CN2016/073024; priority CN 201510071747.0.
- [80] Aspen Plus Brochure — aspentech.com. <https://www.aspentech.com/en/resources/brochure/aspen-plus-brochure>. [Accessed 09-05-2026].
- [81] Aspen Custom Modeler | Quick and Easy | AspenTech — aspentech.com. <https://www.aspentech.com/en/products/engineering/aspen-custom-modeler>. [Accessed 09-05-2026].
- [82] Antoine Rouxhet, Alejandro Morales et Grégoire Léonard : Comparative analysis of recycling loops for sustainable aviation fuel synthesis from CO_2 . F.R.S.-FNRS - Fonds de la Recherche Scientifique, 09 July 2025.

- [83] Antoine Rouxhet, Alejandro Morales et Grégoire Léonard : Study of the base case in a comparative analysis of recycling loops for sustainable aviation fuel synthesis from co₂. F.R.S.-FNRS - Fonds de la Recherche Scientifique, 09 July 2025.
- [84] Antoine Rouxhet et Grégoire Léonard : Towards sustainable aviation fuel from co₂: Energy efficiency assessment of the co₂-to-kerosene pathway. F.R.S.-FNRS - Fonds de la Recherche Scientifique, 05 May 2025.
- [85] Antoine Rouxhet et Grégoire Léonard : The reverse water-gas shift reaction as an intermediate step for synthetic jet fuel production: A reactor sizing study at two different scales. volume 53. F.R.S.-FNRS - Fonds de la Recherche Scientifique, Elsevier, 27 June 2024.
- [86] Antoine Rouxhet et Grégoire Léonard : Using h₂ for co₂ activation on the way towards synthetic fuels. F.R.S.-FNRS - Fonds de la Recherche Scientifique, 20 March 2024.
- [87] Yiping Wu et Shuguang Liu : A suggestion for computing objective function in model calibration. *Ecological Informatics*, 24:107–111, novembre 2014.
- [88] Tikhonov Regularization - an overview | ScienceDirect Topics — sciencedirect.com. <https://www.sciencedirect.com/topics/computer-science/tikhonov-regularization>. [Accessed 09-05-2026].
- [89] minimize(method='Nelder-Mead') SciPy v1.17.0 Manual. <https://docs.scipy.org/doc/scipy/reference/optimize.minimize-neldermead.html>. [Accessed 09-05-2026].
- [90] minimize(method='Powell') SciPy v1.17.0 Manual. <https://docs.scipy.org/doc/scipy/reference/optimize.minimize-powell.html>. [Accessed 09-05-2026].
- [91] C Perego : Experimental methods in catalytic kinetics. *Catalysis Today*, 52(2-3):133–145, 1999.
- [92] D. Toye : Personal communication on reactor sizing, 2026. Personal communication, University of Liège.
- [93] ScienceDirect : Ergun Equation - an overview. <https://www.sciencedirect.com/topics/engineering/ergun-equation>. [Accessed 16-05-2026].
- [94] Luis Vaquerizo et Diego Rego-Fernández : Net-zero sustainable aviation fuel (saf) production via co₂ hydrogenation in low-temperature fischer-tropsch synthesis: Process design and alternatives. *Journal of CO₂ Utilization*, 102:103225, décembre 2025.
- [95] Anthony G. Dixon : Correlations for wall and particle shape effects on fixed bed bulk voidage. *The Canadian Journal of Chemical Engineering*, 66(5):705–708, octobre 1988.
- [96] PubChem : Alumina — pubchem.ncbi.nlm.nih.gov. <https://pubchem.ncbi.nlm.nih.gov/compound/alumina#section=InChI>. [Accessed 04-05-2026].